

Numerical Relativity Injection Infrastructure

Patricia Schmidt,^{1,2,3,*} Ian W. Harry,^{4,†} and Harald P. Pfeiffer^{4,5,‡}

¹*LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, California 91125, USA*

²*Theoretical Astrophysics, Walter Burke Institute for Theoretical Physics,
California Institute of Technology, Pasadena, California 91125, USA*

³*Department of Astrophysics/IMAPP, Radboud University
Nijmegen, P.O. Box 9010, 6525 GL Nijmegen, The Netherlands*

⁴*Max Planck Institute for Gravitational Physics (Albert-Einstein-Institute),
Am Mühlenberg 1, Potsdam-Golm, 14476, Germany*

⁵*Canadian Institute for Theoretical Astrophysics, 60 St. George
Street, University of Toronto, Toronto, ON M5S 3H8, Canada*

This document describes the new Numerical Relativity (NR) injection infrastructure in the LIGO Algorithms Library (LAL), which henceforth allows for the usage of NR waveforms as a discrete waveform approximant in LAL. With this new interface, NR waveforms provided in the described format can directly be used as simulated GW signals (“injections”) for data analyses, which include parameter estimation, searches, hardware injections etc. As opposed to the previous infrastructure, this new interface natively handles sub-dominant modes and waveforms from numerical simulations of precessing binary black holes, making them directly accessible to LIGO analyses. To correctly handle precessing simulations, the new NR injection infrastructure internally transforms the NR data into the coordinate frame convention used in LAL.

I. INTRODUCTION

LIGO reported the first detections of gravitational waves (GW) from merging binary black holes [1–3]. Such coalescing compact binaries are prime sources of GWs for ground-based interferometric GW detectors (see e.g. [4]). Whilst for low-mass systems only the early part of the binary evolution, the inspiral, is accessible to LIGO, for high-mass systems also the later stages, in particular the merger and ringdown of the final black, are visible in LIGO’s sensitivity band.

During the early part of the binary coalescence, the emitted gravitational waveforms are accurately described by analytic post-Newtonian (PN) expansions of the Einstein field equations (see e.g. [5]). To obtain the waveforms through the final stages of the binary coalescence, the full non-linear solutions of the field equations are required, which are provided by Numerical Relativity (NR) [6–8] (see for example [9] for a comprehensive overview).

Numerical Relativity already plays a crucial role in GW data analysis: the construction of waveform models that govern the complete inspiral-merger-ringdown signal (see [10] for a review) depend heavily on NR simulations. Such waveform models [11–14] underpin LIGO’s GW searches, parameter estimation and tests of general relativity (see [1] and references therein). Furthermore, numerical simulations are crucial to determine the remnant black hole’s mass and spin [15], and to investigate systematic biases due to waveform modeling errors [16].

*Electronic address: patricia.schmidt@ligo.org

†Electronic address: ian.harry@ligo.org

‡Electronic address: harald.pfeiffer@ligo.org

In the Advanced detector era, it is very advantageous to be able to directly use NR waveforms in gravitational-wave searches and parameter estimation, to test General Relativity and to assess the systematics of analytic waveform models within a uniform framework. Such use of NR waveforms should be easy and convenient. This is the purpose of the new *Numerical Relativity Injection Infrastructure* described in this document. Once the NR data is provided in a specific format (as described below), this infrastructure allows for the treatment of NR waveforms as a “discrete” waveform approximant, which can seamlessly be called from within the LIGO Algorithm Library (LAL)¹.

In previous efforts, binary-black-hole (BBH) hybrid waveforms constructed by combining a PN inspiral with an NR merger-ringdown waveform, were used in LIGO data analysis and parameter estimation in the NINJA and NINJA-2 projects [17, 18]. However, the previously employed NR modules in LAL require the NR waveforms to be resampled at a uniform time-spacing. In the NINJA framework, the resampling was performed *before* inserting the total mass scale. For the waveforms to be useable at high total mass, the time-spacing in the NR data has to be very small, resulting in very large storage requirements, even if only the dominant harmonics $(\ell, m) = (2, \pm 2)$ were considered.

The new infrastructure described here improves on the earlier approaches in several significant ways. First, data is stored in a highly efficient compressed format [19]; even including a large number of sub-dominant modes (as is now encouraged), storage requirements are lower than just for the dominant modes in the preceding storage format. Second, the compressed NR data are interpolated with one-dimensional spline interpolation *after* the mass scale is inserted. This avoids high-memory operations and further reduces storage requirements and I/O times. Third, the new infrastructure handles projection of the NR data onto arbitrary source-location and detector orientation using the same conventions in LAL as for other waveform families entirely agnostic of the NR code that was used to produce the simulation. The new injection infrastructure is fully implemented in LAL and is intended to supersede previously used NR modules.

The remainder of this technical document is organised as follows: In Sec. II we provide a brief summary of the NR data format and metadata required as input. In Sec. III we describe the basics of the waveform evaluation code and give explicit examples of how the NR waveforms are evaluated in `lalsimulation`. Sec. IV details the frame transformations between the NR frame and the LAL wave-frame. We highlight caveats and desired future improvements in Sec. V.

II. WAVEFORM FORMAT

All data for *one* NR simulation is provided in a single HDF5-file. This file contains:

1. The gravitational waveforms given as spherical-harmonic modes in a spline-compressed format.
2. Metadata describing the simulation, and identifying the origin of the simulation.
3. Optionally, additional information about the dynamics of the black holes.

When multiple NR datasets for the identical physical configuration (e.g. during a convergence test, or for different choices during GW extraction) are provided, then each NR dataset including metadata needs to be stored in a separate .h5 file with a unique name.

A. NR conventions

In Numerical Relativity one solves for the complete space-time of the binary system. For GW data analysis purposes one requires the gravitational-wave strain h far from the source. The relevant numerical quantity is the metric perturbation h_{ij} as computed in the transverse-traceless (TT) gauge.

¹ Available at <https://wiki.ligo.org/DASWG/LALSuite>

There are different ways of computing the metric perturbation from a numerical evolution. The most common methods include the use of the complex Weyl scalar Ψ_4 [20, 21], which is related to the metric perturbation via two time derivatives, or the Regge-Wheeler-Zerilli formalism [22–25], which computes the metric perturbation in the wave-zone as a perturbation of the Schwarzschild spacetime.

In the TT gauge, the metric perturbation has two independent real polarisations, h_+ and h_\times , which can be written as the complex strain

$$h = h_+ - ih_\times \in \mathbb{C}, \quad (1)$$

where $h_+, h_\times \in \mathbb{R}$.

Let $(\hat{e}_x, \hat{e}_y, \hat{e}_z)$ be a Cartesian coordinate system in the *wave-zone*, i.e. the zone far away from the binary where the GWs are extracted. This Cartesian coordinate system is related to the polar coordinates (r, θ, ϕ) by the standard transformation. In this coordinate system, henceforth referred to as the NR frame, the metric perturbation is commonly decomposed into *modes* in a basis of spin-weighted spherical harmonics, $^{-2}Y_{\ell m}$, of spin weight $s = -2$, where the GW propagation direction is the radial unit vector \hat{r} . For any point (θ, ϕ) on the unit sphere, the GW strain takes the form

$$h^{\text{NR}}(t_{\text{GW}}; \theta, \phi) = h_+^{\text{NR}} - ih_\times^{\text{NR}} = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} H_{\ell m}(t_{\text{GW}}) {}^{-2}Y_{\ell m}(\theta, \phi), \quad (2)$$

where $H_{\ell m}(t_{\text{GW}})$ denotes the extracted NR gravitational-wave modes. The modes are given in terms of a retarded time-coordinate t_{GW} . As for any wave, we can also write each mode $H_{\ell m}(t_{\text{GW}})$ as an amplitude $A_{\ell m}(t_{\text{GW}})$ and a phase $\Phi_{\ell m}(t_{\text{GW}})$,

$$H_{\ell m}(t_{\text{GW}}) = A_{\ell m}(t_{\text{GW}}) e^{i\Phi_{\ell m}(t_{\text{GW}})}. \quad (3)$$

The mode amplitude $A_{\ell m}(t)$ is defined as the complex norm of the complex time series $H_{\ell m}(t_{\text{GW}})$, the phase $\Phi_{\ell m}(t_{\text{GW}})$ is the unwrapped argument of the complex time series $H_{\ell m}(t_{\text{GW}})$. For a binary that is orbiting counter-clockwise in the xy -plane of the NR frame (i.e., the orbital angular frequency vector is *parallel* to the z -axis), $\Phi_{2,2}(t_{\text{GW}})$ is a monotonically **decreasing** function.

Following the LAL waveform convention, the time-coordinate t_{GW} in the waveform modes has to be chosen such that the peak of the waveform occurs at $t_{\text{GW,peak}} \equiv 0$, where the ‘‘peak’’ of the waveform is defined by

$$h_{\text{peak}} := \max \left(\sum_{\ell, m} A_{\ell m}(t_{\text{GW}})^2 \right). \quad (4)$$

Note that the time coordinate t_{GW} in the wave-zone is not identical to time coordinate t_{NR} used by the numerical relativity simulation in the strong-field regime, the latter is used below to give information about the dynamics of the black holes. There is no unambiguous identification of t_{GW} with t , since they are defined in different regions of the space-time (far-zone vs. near-zone). The normalization of t_{GW} is given by Eq. (4); we describe on page 5 how to choose t .

B. Amplitude and phase spline compression

Gravitational waveforms for LIGO data analysis purposes require uniform sampling in time for a given sampling frequency. NR datasets, however, are commonly not uniformly sampled and if they are, the sampling interval Δt_{GW} may not necessarily correspond to the one required by data analysis tools. It is therefore unavoidable to interpolate the NR data to the desired sampling rate. Whilst the NR data could simply be interpolated as they are, we choose to reduce the data by performing one-dimensional spline

compression on the NR data [19]. This is a particular advantage for long simulations or hybrid data, but also significantly reduces the storage and I/O for pure NR data.

The one-dimensional spline compression is performed separately for each **mode amplitude** $A_{\ell m}(t_{\text{GW}})$ and **mode phase** $\Phi_{\ell m}(t_{\text{GW}})$, which are already time-shifted such that $t_{\text{GW}} = 0$ corresponds to the peak of the waveform. Pure NR data without an inspiral need to have the initial *junk radiation* removed *before* the spline interpolants are constructed. We refer to this very first data point stored in the time series after removing the initial junk as the *beginning* of the waveform. No tapering should be applied at the beginning or end of the waveform data.

The routine employed to compress and interpolate the data is the *reduced-order spline interpolation* presented in [19]. It uses a greedy algorithm that selects the near-optimal points to construct a univariate spline interpolant with a specified global accuracy and polynomial degree. By default, the interpolants are constructed using fifth degree polynomials and a tolerance of 10^{-6} , i.e., if the spline is evaluated at the original discrete NR times t_{GW} , the original NR values for mode amplitude and phase are recovered with an error equal to or smaller than the specified tolerance. In addition, a median error can be associated with any predicted value, i.e., a value not contained in the original NR data. For a detailed description of this method and the accuracy of the obtained interpolants, we refer the reader to [19]. The spline compression is conveniently performed using the publicly available Python package `romSpline` by Chad R. Galley [26]. The spline interpolants for each amplitude and phase are obtained via `romSpline` as follows:

```
import romSpline
spline = romSpline.ReducedOrderSpline(tGW, Aℓ,m(tGW) or Φℓ,m(tGW),
                                       verbose=False)

spline.write('filename.h5')
```

The output of `spline.write` contains all information needed for subsequent interpolation of the respective input time-series data and is composed of five datasets: (deg, tol, X, Y, errors)².

The spline interpolants for each (ℓ, m) -mode for a single NR simulation are stored as individual H5-groups in the *final HDF5 output file* under the mandatory group-names given below and summarized in Sec. IID:

Group-name	Type	Description
amp_l#1_m#2	romSpline-group	amplitude $A_{\ell m}(t_{\text{GW}})$
phase_l#1_m#2	romSpline-group	phase $\Phi_{\ell m}(t_{\text{GW}})$

Here (#1, #2) are placeholders for (ℓ, m) , e.g., for $(\ell = 2, m = -2)$ the group naming convention is `phase_l2_m-2` and `amp_l2_m-2`.

C. Metadata stored as HDF5-attributes

The metadata format is adapted from the original NINJA-2 metadata format [27]. The metadata are described in the following list, and are stored as *attributes* of the final HDF5 file for each NR simulation. Often, more extensive metadata are available for a NR simulation. We recommend that such additional metadata is included as a free-format text dataset in the 'auxiliary-info' H5-group (see below).

Required time-dependent metadata (see list below) have to contain the values corresponding to the first entry in the stored time series. For pure NR data, since the junk radiation has to be removed, these are

² The latest version of `romSpline` allows to pass a group descriptor, so that the spline-data can be written directly into the appropriate group of the final output file.

not the initial data of the simulation. We refer to this very first data point stored in the time series as the *beginning* of the waveform data, subsequently indicated by t_{begin} .

Vectors must be represented by their Cartesian components in the NR frame (see Sec. II A). This coordinate system is consistent with the one used in the decomposition in spherical harmonics (e.g. the polar axis of the spherical harmonics must point along the +z-axis). No requirements are placed on the orientation of the NR frame. For instance, the orbital unit separation vector ' \hat{n} ' and the Newtonian unit orbital angular momentum ' \hat{L} ' need not point along specific coordinate axes. Note also that the spin-fields in the metadata ' $\text{spin}\{1,2\}\{x,y,z\}$ ' are given in the NR frame, and *not* in the LAL frame. The rotation of the GW modes into the LAL frame is taken care of by the NR injection infrastructure, following the conventions described in Sec. IV.

Note: The time-coordinate of spins, positions, masses and their respective time-series will be associated with the apparent horizon, e.g. the NR coordinate time t . Such a time cannot unambiguously be identified with the retarded time t_{GW} of the waveform modes. However, the submitter is asked to make a **reasonable effort to have the same numerical values of t and t_{GW} to correspond to the same portion of the waveform**, e.g. through choosing $t_{\text{GW}} = t - R_{\text{extraction}}^*$, where $R_{\text{extraction}}^*$ is the tortoise radius of the extraction sphere [28, 29].

<i>Attribute-name</i>	<i>Type</i>	<i>Description</i>
Format	integer	indicates what data are supplied. Must be 1, 2, 3
type	string	keyword description of the simulation as requested by the LIGO Open Science Center (LOSC). Required value: NRinjection
name	string	short identifier of the simulation, e.g., SXS:BBH:0019
alternative-names	string	comma-separated list of user-specifiable alternative names. These names can be longer, more descriptive, and/or include what specific series of simulations this configuration belongs to.
NR-group	string	name of the NR group that carried out the simulation
NR-code	string	name of the NR code that was used to carry out the simulation
modification-date	string	date when this .h5 file was last updated. Format 'YYYY-MM-DD'
point-of-contact-email	string	contact person for questions
simulation-type	string	keyword description of the spin configuration. Allowed values are: aligned-spins, non-spinning, precessing
INSPIRE-bibtex-keys	string	comma-separated list of INSPIRE bibtex keys that should be cited when this waveform is used (1-3 publications)
license	string	allowed values are: LVC-internal, public
Lmax	integer	the maximum ℓ -value for which $h_{\ell m}$ -modes are supplied (all modes with $\ell \leq L_{\text{max}}$, $-\ell \leq m \leq +\ell$ have to be provided)
NR-techniques	string	attempts to summarize major elements of the NR simulation. A comma-separated list of one element in each of the following categories

continued

<i>Attribute-name</i>	<i>Type</i>	<i>Description</i>
		Category 1: Puncture-ID, Quasi-Equilibrium-ID Category 2: BSSN, GH, Z4c Category 3: RWZ-h, Psi4-integrated Category 4: Finite-Radius-Waveform, CCE-Waveform, Extrapolated-Waveform Category 5: ApproxKillingVector-Spin, CoordinateRotation-Spin Category 6: Christodoulou-Mass Example: 'Quasi-Equilibrium-ID, GH, RWZ-h, Extrapolated-Waveform, ApproxKillingVector-Spin, Christodoulou-Mass' (Note: This list is extensible. If your code does not fit the given choices, contact the authors.)
files-in-error-series	string	a comma-separated list of .h5 files (including the present one) that combined form an error series for the binary configuration, e.g. different numerical resolutions. Set to ' ' if no error-series for this configuration exists.
comparable-simulation	string	one other .h5 file that (a) has an error-series and (b) is numerically “comparable” to the present one, i.e. an error-analysis that is performed on 'comparable-simulation' is expected to carry over to this waveform. Set to ' ' if an error-series is provided.
production-run	integer	allowed values are 1 and 0. If 1, this is the highest quality member of the error-series and should be used for analyses. If 0, this is a lower-quality member of the error-series and should not be used for general analyses.
object1	string	keyword description to identify the object type. Allowed values are: BH, NS
object2	string	keyword description to identify the object type. Allowed values are: BH, NS
mass1	float	mass of the more massive object at t_{begin} ; if both objects are BH, the unit of mass is arbitrary. If at least one object is a NS, then the unit is solar mass M_{\odot} .
mass2	float	mass of the lighter object at t_{begin} ; if both objects are BH, the unit of mass is arbitrary. If at least one object is a NS, then the unit is solar mass M_{\odot} .
eta	float	the symmetric mass ratio of the simulation at t_{begin} .
f_lower_at_1MSUN	float	frequency of the (2,2)-mode in Hz at the beginning of the waveform scaled to $1M_{\odot}$
spin1x	float	x-component of the dimensionless spin vector $\vec{\chi}_1(t_{\text{begin}})$ in NR frame
spin1y	float	y-component of the dimensionless spin vector $\vec{\chi}_1(t_{\text{begin}})$ in NR frame
spin1z	float	z-component of the dimensionless spin vector $\vec{\chi}_1(t_{\text{begin}})$ in NR frame
spin2x	float	x-component of the dimensionless spin vector $\vec{\chi}_2(t_{\text{begin}})$ in NR frame
spin2y	float	y-component of the dimensionless spin vector $\vec{\chi}_2(t_{\text{begin}})$ in NR frame
spin2z	float	z-component of the dimensionless spin vector $\vec{\chi}_2(t_{\text{begin}})$ in NR frame
LNhatx	float	x-component of the Newtonian orbital angular momentum unit vector $\hat{L}_N(t_{\text{begin}})$ in NR frame
LNhaty	float	y-component of the Newtonian orbital angular momentum unit vector $\hat{L}_N(t_{\text{begin}})$ in NR frame
LNhatz	float	z-component of the Newtonian orbital angular momentum unit vector $\hat{L}_N(t_{\text{begin}})$ in NR frame
nhatx	float	x-components of the orbital separation unit vector $\hat{n}(t_{\text{begin}})$ given by Eq. (15) in NR frame

continued

<i>Attribute-name</i>	<i>Type</i>	<i>Description</i>
nhaty	float	y-components of the orbital separation unit vector $\hat{n}(t_{\text{begin}})$ given by Eq. (15) in NR frame
nhatz	float	z-components of the orbital separation unit vector $\hat{n}(t_{\text{begin}})$ given by Eq. (15) in NR frame
Omega	float	dimensionless orbital frequency $M\Omega(t_{\text{begin}})$
eccentricity	float	estimated eccentricity of the simulation at t_{begin}
mean_anomaly	float	estimated mean anomaly (cf. Eq. 21) at t_{begin} . For eccentricity $\leq 10^{-3}$, it is allowed to set mean_anomaly to 0. If the mean anomaly has not been computed, set mean_anomaly to -1.
PN_approximant	string	<i>Only present for PN-NR hybrid waveforms:</i> identifier of the inspiral approximant

D. Data stored as HDF5-groups and datasets

The following groups are required inside the .h5 file that represents a simulation. Some groups are optional and only need to be given if Format=2 or Format=3. All time-series represent ROM-compressed data obtained via `romSpline` in the same way as the amplitudes and phases (see Sec. II B for details).

<i>Group-name</i>	<i>Type</i>	<i>Description</i>
auxiliary-info	H5-group	Contains anything that the submitter finds helpful to identify, document and repeat the run
NRtimes	H5-dataset	<i>optional but highly recommended:</i> 1-d array of the discrete times t_{GW} that formed the input-times into the romSpline compression of <code>amp_l#1_m#2</code> and <code>phase_#1_m#2</code>
GW modes: <i>Two groups for each (ℓ, m), $2 \leq \ell \leq L_{\text{max}}$, $-\ell \leq m \leq +\ell$</i>		
amp_l#1_m#2	romSpline-group	amplitude $A_{\ell m}(t_{\text{GW}})$
phase_l#1_m#2	romSpline-group	phase $\Phi_{\ell m}(t_{\text{GW}})$
if Format ≥ 2, also specify the following time-series:		
mass1-vs-time	romSpline-group	mass of the more massive object
mass2-vs-time	romSpline-group	mass of the less massive object
spin1x-vs-time	romSpline-group	x-component of the dimensionless spin $\vec{\chi}_1(t)$
spin1y-vs-time	romSpline-group	y-component of the dimensionless spin $\vec{\chi}_1(t)$
spin1z-vs-time	romSpline-group	z-component of the dimensionless spin $\vec{\chi}_1(t)$
spin2x-vs-time	romSpline-group	x-components of the dimensionless spin $\vec{\chi}_2(t)$
spin2y-vs-time	romSpline-group	y-components of the dimensionless spin $\vec{\chi}_2(t)$
spin2z-vs-time	romSpline-group	z-components of the dimensionless spin $\vec{\chi}_2(t)$
position1x-vs-time	romSpline-group	x-component of the center $\vec{c}_1(t)$ of object 1
position1y-vs-time	romSpline-group	y-component of the center $\vec{c}_1(t)$ of object 1
position1z-vs-time	romSpline-group	z-components of the center $\vec{c}_1(t)$ of object 1

continued

<i>Group-name</i>	<i>Type</i>	<i>Description</i>
position2x-vs-time	romSpline-group	x-component of the center $\vec{c}_2(t)$ of object2
position2y-vs-time	romSpline-group	y-component of the center $\vec{c}_2(t)$ of object2
position2z-vs-time	romSpline-group	z-component of the center $\vec{c}_2(t)$ of object2
LNhatx-vs-time	romSpline-group	x-component of Newtonian angular momentum direction $\hat{L}_N(t)$
LNhaty-vs-time	romSpline-group	y-component of Newtonian angular momentum direction $\hat{L}_N(t)$
LNhatz-vs-time	romSpline-group	z-component of Newtonian angular momentum direction $\hat{L}_N(t)$
Omega-vs-time	romSpline-group	dimensionless orbital frequency $M\Omega(t)$
If Format ≥ 3, also specify the following time-series		
remnant-mass-vs-time	romSpline-group	remnant mass
remnant-spinx-vs-time	romSpline-group	x-component of the dimensionless spin of the remnant
remnant-spiny-vs-time	romSpline-group	y-component of the dimensionless spin of the remnant
remnant-spinz-vs-time	romSpline-group	z-component of the dimensionless spin of the remnant
remnant-positionx-vs-time	romSpline-group	x-component of the center of the remnant
remnant-positiony-vs-time	romSpline-group	y-component of the center of the remnant
remnant-positionz-vs-time	romSpline-group	z-component of the center of the remnant

III. NR WAVEFORM EVALUATION IN LAL

Once the HDF5 file has been provided, the NR waveforms can be evaluated through the standard waveform interfaces `ChooseTDWaveform` in LAL. The approximant name is “NR_hdf5”. The spline data are read from file and evaluated for the desired extrinsic parameters, total mass and starting frequency. Since some intrinsic parameters of an NR simulations are fixed (e.g. mass-ratio and dimensionless spins), internal checks on the mass ratio and the spin components are performed to guarantee the consistency between the values passed in the waveform generation call and metadata values. Note: The waveform generator requires the input spin values to be defined as given by Eqn. (43). In general, these are different to the values of the spin metadata and need to be computed using the metadata for the spins, the orbital angular momentum and the orbital separation. The function `SimInspiralNRWaveformGetSpinsFromHDF5File` returns the spins the required convention.

For a given starting frequency and total mass, a time array is allocated based on an estimate of the waveform length. We use the LAL-function `SimIMRSEOBNRv2ChirpTimeSingleSpin` to estimate the waveform length with an additional leverage of 10%. If the NR waveforms are not long enough for a given total mass and starting frequency, the generation is aborted and an error is generated. From the estimated length and the desired sampling rate, the discrete time series for the spline evaluation is determined.

To construct the NR GW polarisations h_+ and h_\times in the LAL wave-frame, first the splines for each NR amplitude and phase are first evaluated at the required sampling times and convolved with the spin-weighted spherical harmonics. Finally, the NR polarizations are transformed into the LAL wave-frame following Eqs. (36). Note that all (ℓ, m) -modes present in the HDF5 file are used to compute the two polarisations.

The compressed NR data files do not store the splines themselves, but the X-data, Y-data, errors, the polynomial degree etc. A regular GSL interpolation is used to construct the splines from the HDF5 file. A comparison with the scipy function `UnivariateSpline` found that the mismatch between waveforms reconstructed using the two different interpolators was less than 10^{-7} . This is consistent with the level

of disagreement expected due to the different numerical interpolation routines. Fig. 1 shows an example comparison between NR waveforms obtained using the two different interpolation routines. The source code can be found in `lalsuite/lalsimulation/src/LALSimIMRNRWaveforms.c`.

A. Examples

There are a variety of different ways to evaluate NR waveforms using LIGO data analysis software. Here, we give an explicit example using Python and the SWIG-wrapped version of `lalsimulation`. The only difference between this and generating a waveform using any other waveform model is that the path to the HDF5 file must be provided explicitly, as illustrated.

Example using `lalsimulation` through SWIG:

```
import lal
import lalsimulation as lalsim
# Compute spins in the LAL frame
s1x, s1y, s1z, s2x, s2y, s2z =
lalsim.SimInspirialNRWaveformGetSpinsFromHDF5File('/PATH/TO/H5File')
# Create a dictionary and pass /PATH/TO/H5File
params = lal.CreateDict()
lalsim.SimInspirialWaveformParamsInsertNumRelData(params, '/PATH/TO/H5File')
# Generate GW polarisations
hp, hc = lalsim.SimInspirialChooseTDWaveform(mass1 * MSUN_SI,
      mass2 * MSUN_SI,
      s1x, s1y, s1z,
      s2x, s2y, s2z,
      distance, inclination,
      phiRef, $pi/2$,
      0., 0., deltaT,
      fStart, fRef,
      params, approximant=lalsim.NR_hdf5)
```

Fig. 1 shows the two waveform polarizations h_+ and h_\times for the precessing binary black hole hole simulated in case SXS:BBH:0006 from the publicly available SXS catalogue [31] for a total mass of $50 M_\odot$ and an inclination of $\pi/3$. The dimensionless spins for this simulation in the LAL frame are $\vec{\chi}_1 = (-0.05, -0.27, -0.16)$ and $\vec{\chi}_2 = (-0.02, -0.11, -0.10)$ and the component masses are $m_1 = 28.68$ and $m_2 = 21.32$. The waveform is generated from its beginning, corresponding to the starting frequency of $f_{\text{Start}}=18.76\text{Hz}$. Further parameters are: $\text{distance}=100\text{Mpc}$, $\text{deltaT}=1.0/16384$, $\text{phiRef}=0$ and $\text{fRef}=f_{\text{Start}}$.

IV. LAL COORDINATE FRAMES FOR PRECESSING BINARIES AND NR INJECTIONS

1. Executive summary: changes of conventions

It has recently become apparent that certain waveform conventions in LAL are not ideal to specify precessing binaries: (i) The phase-angle `phiRef` couples the specification of the line-of-sight to Earth with the specification of spin components `s1x`, `s1y`, `s1z`, `s2x`, `s2y`, `s2z`. To compute waveforms for the *identical* compact binary viewed from different directions, one may have to specify different values for the spin-components `s1x`, `s1y`, `s2x`, `s2y`. (ii) Several semi-analytic waveform models do not conform to this convention, already following the convention detailed below.

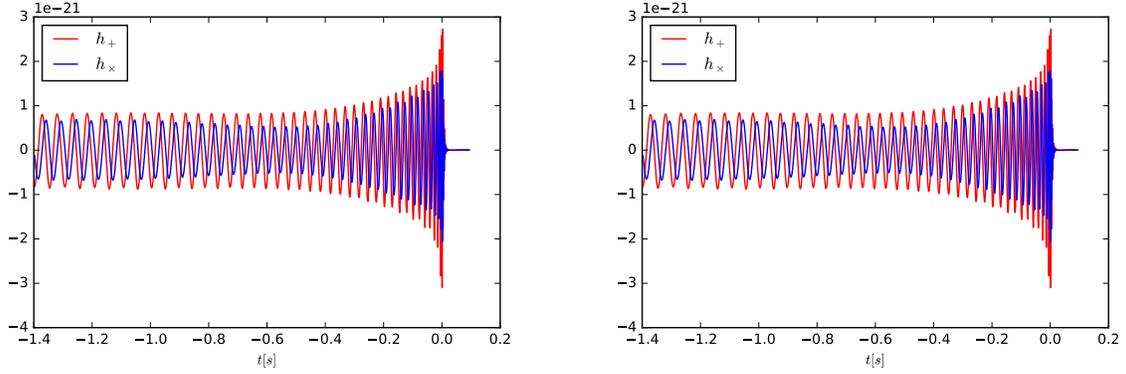


FIG. 1: The waveform polarizations h_+ (red) and h_\times (blue) of the publicly available SXS waveform SXS:BBH:0006 generated via the waveform interface in `lalsimulation` (left panel) and constructed using `scipy`'s `UnivariateSpline` function (right panel).

Concretely, the spin-components are now specified in a geometric way based on the angular momentum \hat{L} and the line connecting object 2 to object 1, \hat{n} :

$$S1_x \stackrel{\text{ref}}{=} \vec{\chi}_1 \cdot \hat{n}, \quad (5)$$

$$S1_y \stackrel{\text{ref}}{=} \vec{\chi}_1 \cdot (\hat{L} \times \hat{n}), \quad (6)$$

$$S1_z \stackrel{\text{ref}}{=} \vec{\chi}_1 \cdot \hat{L}, \quad (7)$$

(and similarly for the second object). The symbol $\stackrel{\text{ref}}{=}$ indicates that equality only holds at a reference time as spins generically precess during an inspiral [32, 33].

Furthermore, it is suggested to specify orientation of eccentric orbits through the angles

$$\text{phiRef} = \Phi \stackrel{\text{ref}}{=} \angle(\hat{n}, \text{line-of-ascending-node}), \quad (8)$$

$$\text{mean_anomaly} = \delta. \quad (9)$$

2. Motivation & benefits of new conventions

Generally waveform modeling requires at least two coordinate systems, a “source-frame” in which it is convenient to specify properties of the source of gravitational waves, and a “wave-frame” which is adopted to wave-propagation to GW detectors on Earth. Furthermore, NR data are specified in whatever coordinates are employed during the numerical evolution, generally resulting in a third coordinate system. This section defines coordinate frames for use in LAL. Specifically, we achieve:

1. Identification of a set of intrinsic parameters that fully describe the dynamics of a binary on an eccentric orbit, defined solely in terms of the source-frame (i.e. independent of the wave-frame).
2. Identification of three angles that describe the transformation between source- and wave-frame, which are independent of the intrinsic parameters.
3. Identification of parameters describing the orbital phase and periapsis location, which have a convenient circular-orbit limit: As the eccentricity tends to zero, one of the two phase-parameters reduces to the standard orbital phase for circular orbits, whereas the other becomes irrelevant.
4. Identities that relate the basis-vectors in the source-frame to the wave-frame (and vice versa). These identities are written in vectorial form and are valid in any coordinate system.

The orthogonal decomposition into intrinsic and extrinsic parameters (points 1 and 2) allows to change the direction at which a binary is viewed (i.e. the wave-frame), without having to adjust the parameters that determine the intrinsic dynamics. Point 3 prepares the ground for easy extension to eccentric waveforms. point 4 is of particular relevance when translating NR data into LAL-conventions: Evaluating the vector identities in the NR-coordinate system, yields immediately the relation between NR coordinates and wave-frame.

3. NR waveform injections

NR data are assumed to be supplied in the data-format defined in Sec. II and III, and so it needs to be transformed into the LAL wave-frame. For a given reference time, an entire NR data-set could in principle be transformed into the LAL frame by suitably transforming each $H_{\ell m}$ -mode. Applying such a coordinate transformation on the NR-data as a pre-processing step before using it for injections suffers from two disadvantages: First, the person who prepares the NR-data for LAL-use must perform the rotation, and must do so correctly. Since NR data is prepared separately by several NR groups, this opens the possibility of introducing errors in this step. Secondly, for precessing systems the orbital angular momentum \hat{L} precesses. Therefore, pre-rotating the NR-data locks in the reference point, and one would need to generate different pre-transformed NR-data for different reference points. Since NR data is in geometric units, the NR-data would have to be separately transformed whenever the total mass of an injection changes.

To avoid both disadvantages, it is proposed to leave the NR-data in its original frame (see Sec. II). Instead, the transformations between the LAL-convention and the NR-data are applied *at use* during the call to the waveform evaluation function in LAL (cf. Sec. III).

This section develops the necessary transformations to implement this technique. No assumptions are made on the NR-coordinate system in order to allow for precession when NR-orbital angular momentum is typically *not* along the z-axis of the NR-coordinate system. The derived formulae also allow to choose an arbitrary reference point for the NR waveform, as long as orbital angular momentum and vector connecting the two compact objects are known at that point.

A. Coordinate frames

The frames defined here differ somewhat from the preceding LAL conventions as detailed in [34]. The differences are needed to achieve the separation between intrinsic and extrinsic parameters. In the old conventions, the orbital phase ϕ was also used as part of the rotation parameters that define the rotation between source- and wave-frame. Therefore, to “look” at the same binary from different angles used to require a suitable change in the spin-components tangential to the orbital plane.

1. NR frame ($\hat{e}_x, \hat{e}_y, \hat{e}_z$)

This is a generic coordinate system without any regards to the concrete binary motion. Generic coordinate systems occur in numerical relativity, where the coordinates are chosen through some gauge-conditions, and the binary is evolving from some initial data. At some later time therefore, the coordinates will not have any particular, controlled properties. The Cartesian basis-vectors are denoted

$$\hat{e}_x, \hat{e}_y, \hat{e}_z. \tag{10}$$

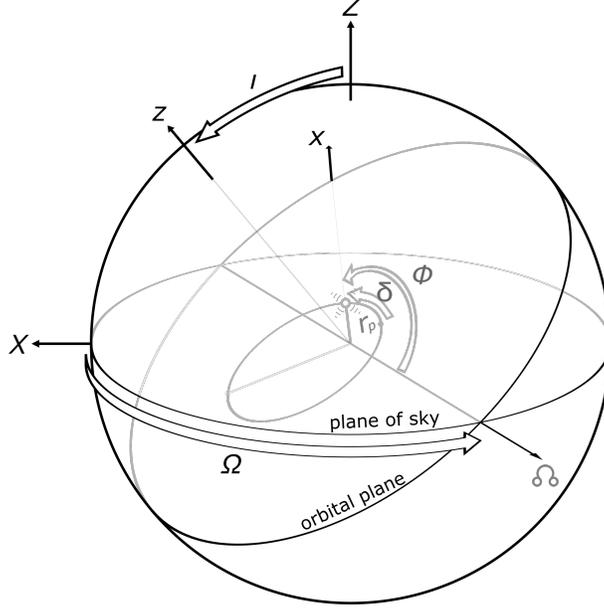


FIG. 2: Coordinate frames. (x,y,z) denotes the source-frame and (X,Y,Z) the wave-frame. For the recommended default $\Omega = \pi/2$, the line of ascending nodes agrees with \hat{Y} i.e. the rotation by the inclination angle i is about the \hat{Y} -axis.

From these, spherical basis-vectors can be computed as

$$\hat{r} = \cos \phi \sin \theta \hat{e}_x + \sin \phi \sin \theta \hat{e}_y + \cos \theta \hat{e}_z \quad (11a)$$

$$\hat{\theta} = \cos \phi \cos \theta \hat{e}_x + \sin \phi \cos \theta \hat{e}_y - \sin \theta \hat{e}_z \quad (11b)$$

$$\hat{\phi} = -\sin \phi \hat{e}_x + \cos \phi \hat{e}_y. \quad (11c)$$

NR data are assumed to be represented by spherical-harmonic modes of the NR coordinates, as previously defined in the NINJA data-formats document [27]:

$$h_+^{\text{NR}}(t_{\text{GW}}) = \frac{1}{2} (\hat{\theta}_i \hat{\theta}_j - \hat{\phi}_i \hat{\phi}_j) h^{ij}(t_{\text{GW}}) \quad (12a)$$

$$h_{\times}^{\text{NR}}(t_{\text{GW}}) = \frac{1}{2} (\hat{\theta}_i \hat{\phi}_j + \hat{\phi}_i \hat{\theta}_j) h^{ij}(t_{\text{GW}}), \quad (12b)$$

and

$$\sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} H_{\ell m}(t_{\text{GW}}) {}^{-2}Y_{\ell m}(\theta, \phi) = \frac{r}{M} \left(h_+^{\text{NR}}(t_{\text{GW}}) - i h_{\times}^{\text{NR}}(t_{\text{GW}}) \right). \quad (13)$$

As explained in Sec. II, the data-files are assumed to contain a compressed time-series of amplitudes and phases per Eq. (3). Given an emission direction (θ, ϕ) , Eq. (13) yields the GW modes $h_+^{\text{NR}}(t_{\text{GW}})$ and $h_{\times}^{\text{NR}}(t_{\text{GW}})$, according to the convention Eqs. (12).

The gravitational wave data are given in a time-coordinate t_{GW} of observers at large distance r .

Let us now turn to a description of the dynamics of the two bodies. Numerical relativity defines a variety of vectors, the combination of which defines the instantaneous state of the two bodies. These are: The dimensionless spin vectors of the two bodies,

$$\vec{\chi}_1(t), \vec{\chi}_2(t); \quad (14)$$

the direction from body 2 to body 1,

$$\hat{n}(t) := \frac{\vec{c}_1 - \vec{c}_2}{\|\vec{c}_1 - \vec{c}_2\|}, \quad (15)$$

where \vec{c}_i is the coordinate centre of the horizon of the i -th body; the direction of the Newtonian orbital angular momentum,

$$\hat{L}_N(t). \quad (16)$$

We do **not** make any assumption about the relation of the NR orbital angular momentum vector \hat{L} relative to the NR coordinates³. One can further define an orbital frequency

$$\vec{\Omega}_{\text{orb}}(t) = \hat{n}(t) \times \frac{d\hat{n}(t)}{dt}. \quad (17)$$

Equations (14)–(17) are defined in the strong-field region near the black holes, and are given as a function of the time-coordinate t employed by the NR simulation in the strong field regime. The time-coordinates t_{GW} and t are defined at different regions of the space-time (far-zone vs. near-zone), and their preferred relative alignment is discussed on page II C. Any relation between t_{GW} and t , however, is ambiguous. Often NR simulations employ an approximation of “retarded time”, i.e.

$$t_{\text{GW}} = t|_{\text{retarded}}. \quad (18)$$

However, in a dynamical space-time with black holes, two difficulties arise: First, the horizons are causally disconnected from future null infinity, so there are no outgoing null-rays that connect the horizons to the wave-zone. Secondly, when integrating null-rays “slightly outside” the horizons, the time-delay to infinity will depend on precisely where the integration was started, as well as on the initial direction of the null-ray. Therefore, any relation between t_{GW} and t should be viewed as approximate. One should further assume that different NR groups may use different definitions of t_{GW} in the data they compute⁴.

2. LAL Source-Frame ($\hat{x}, \hat{y}, \hat{z}$)

The LAL source-frame ($\hat{x}, \hat{y}, \hat{z}$) is defined as follows:

1. The \hat{z} -axis points along the orbital angular momentum of the binary,

$$\hat{z} \stackrel{\text{ref}}{=} \hat{L}. \quad (19)$$

2. The \hat{x} -axis points along the vector \hat{n} pointing from the second to the first body,

$$\hat{x} \stackrel{\text{ref}}{=} \hat{n}. \quad (20)$$

3. The third vector $\hat{y} = \hat{z} \times \hat{x}$ completes the triad.

³ The customary choice of many NR groups is to *start* NR simulations with $\hat{L} = \hat{e}_z$. Because of junk-radiation and precession effects, $\hat{L}(t)$ will deviate from \hat{e}_z . This deviation generally will be very small for aligned-spin BBH systems, but may become significant for precessing systems.

⁴ S_{PEC}, for example, reports GW-waveforms extracted at *finite-radius* in terms of the NR coordinate, $t_{\text{GW}} = t$, whereas extrapolated waveforms are reported using a retarded time-coordinate with a correction of the rate of flow of time, see Eqs. (7), (14a) and (14b) of Ref. [29].

The symbol $\stackrel{\text{ref}}{=}$ indicates that the respective equation is only required at the reference epoch. The reference epoch can be unambiguously specified by a reference-time t_{ref} . One can also specify a reference *orbital* frequency $\Omega_{\text{orb,ref}}$, and infer t_{ref} via Eq. (17), $\Omega_{\text{orb}}(t_{\text{ref}}) = \Omega_{\text{orb,ref}}$. If the reference epoch is desired to be specified in terms of a *gravitational wave*-frequency, then the gravitational wave time t_{GW} needs to be related to the time-coordinate of the black hole dynamics, t . As discussed in the context of Eq. (18), such an identification is ambiguous and holds only approximately.

Equations (19) and (20) define the source-frame at the reference epoch only. They are chosen such that the spin-components (S_{x1}, S_{y1}, S_{z1}) and (S_{2x}, S_{2y}, S_{2z}) have coordinate-invariant meaning: S_{x1} is the projection of $\vec{\chi}_1$ onto \hat{n} , S_{z1} is the projection of $\vec{\chi}_1$ onto \hat{L} , etc.

The source-frame does not rotate as the binary evolves. Specifically, the rotation between source-frame and wave-frame described below is constant in time.

A different reference epoch would lead to a different source-frame, related by some rotation. If the reference epoch is shifted by a small amount (comparable to the orbital time-scale), \hat{x} and \hat{y} would rotate with the binary. On the precession time-scale \hat{z} would change.

The source frame has no deep intrinsic, geometric significance. It is merely a vehicle to describe the spin-projections onto intrinsic geometric directions, and the basis-vectors ($\hat{x}, \hat{y}, \hat{z}$) will be convenient when writing down the transformation to the wave-frame.

3. Intrinsic parameters of a binary

Given a reference epoch, a binary is specified by the following ten numbers:

- Two masses m_1, m_2 .
- Two spin-vectors $\vec{\chi}_1, \vec{\chi}_2$, specified through the projections of the spin-vectors onto \hat{L}, \hat{n} and the third basis-vector (at reference time).
- Eccentricity e .
- Mean anomaly δ . The mean anomaly is defined in terms of the current time t , the time of last periapsis passage T_{previous} , and the time of next periapsis passage T_{next} :

$$\delta = 2\pi \frac{t - T_{\text{previous}}}{T_{\text{next}} - T_{\text{previous}}}. \quad (21)$$

Eccentricity should be defined somehow “near” the reference epoch. A precise definition (if possible at all) is left to the future. For low eccentricity orbits, there is a competition between radiation reaction driven inspiral (which results in a slightly negative average radial velocity), and the oscillatory radial motion due to eccentricity. For sufficiently small eccentricity, minima in separation may no longer exist. If mean anomaly is still needed despite the quite low eccentricity in such cases, one will have to define periapsis as minimum of separation compared to a fiducial smooth inspiral trajectory.

4. Wave-Frame ($\hat{X}, \hat{Y}, \hat{Z}$)

The wave-frame is adopted to the direction of the observer (i.e. Earth), such that its \hat{Z} -axis points toward the observer. \hat{X} and \hat{Y} represent basis-vectors orthogonal to the line-of-sight, i.e. they span the plane of the sky. The wave-frame is completely specified by three angles:

1. The angle Φ between line of ascending node and \hat{n} (at the reference time).
2. The inclination ι , i.e. the angle between orbital angular momentum $\hat{L} = \hat{z}$ and the line-of-sight \hat{Z} .
3. The angle Ω between the \hat{X} -axis and the line of the ascending node (at the reference time).

These three angles happen to be the Euler angles of the rotation from the wave-frame to the source-frame. In the LAL wave-frame, the gravitational wave modes are defined as

$$h_+^{\text{W}} = \frac{1}{2} (\hat{X}_i \hat{X}_j - \hat{Y}_i \hat{Y}_j) h^{ij}, \quad (22a)$$

$$h_{\times}^{\text{W}} = \frac{1}{2} (\hat{X}_i \hat{Y}_j + \hat{Y}_i \hat{X}_j) h^{ij}, \quad (22b)$$

where the superscript 'W' indicates the LAL wave-frame. The angle Ω rotates \hat{X} and \hat{Y} into each other. By definition of Eq. (22), this merely rotates the polarization of the GW modes, and so Ω is fully degenerate with the GW polarization.

Note that the phase-angle Φ is specified without regard to the location of periapsis of the binary (this differs from the previous LAL convention). This new definition decouples the specification of periapsis location and the specification of orbital phase, and avoids the ambiguity that would arise in the zero-eccentricity limit of definitions involving periapsis. Indeed, for fixed Φ and fixed δ , as eccentricity approaches zero, the waveforms will approach the identical circular form, independent of the value of δ .

B. Relationship between frames

1. From source-frame to wave-frame

The relation between source-frame and wave-frame can be easily derived following the three rotations that rotate one frame into the other: Beginning in the source-frame, $\hat{x}, \hat{y}, \hat{z}$, we first apply the rotation by Φ around \hat{z} , which yields a frame with basis-vectors $\hat{p}, \hat{q}, \hat{z}$:

$$\hat{p} = \cos \Phi \hat{x} - \sin \Phi \hat{y}, \quad (23a)$$

$$\hat{q} = \sin \Phi \hat{x} + \cos \Phi \hat{y}, \quad (23b)$$

$$\hat{z} = \hat{z}. \quad (23c)$$

The inverse rotation is

$$\hat{x} = \cos \Phi \hat{p} + \sin \Phi \hat{q}, \quad (24a)$$

$$\hat{y} = -\sin \Phi \hat{p} + \cos \Phi \hat{q}, \quad (24b)$$

$$\hat{z} = \hat{z}. \quad (24c)$$

The vectors \hat{p}, \hat{q} form an orthonormal basis of the \hat{x} - \hat{y} plane, such that \hat{p} points in the direction of ascending node.

Next, we rotate around the line of the ascending node by the inclination ι , resulting in basis-vectors

$$\hat{P} = \hat{p}, \quad (25a)$$

$$\hat{Q} = \cos \iota \hat{q} - \sin \iota \hat{z}, \quad (25b)$$

$$\hat{Z} = \sin \iota \hat{q} + \cos \iota \hat{z}. \quad (25c)$$

\hat{P} and \hat{Q} form an orthonormal basis of the \hat{X} - \hat{Y} -plane with \hat{P} pointing in the direction of the ascending node. The inverse rotation is

$$\hat{p} = \hat{P}, \quad (26a)$$

$$\hat{q} = \cos \iota \hat{Q} + \sin \iota \hat{Z}, \quad (26b)$$

$$\hat{z} = -\sin \iota \hat{Q} + \cos \iota \hat{Z}. \quad (26c)$$

The final rotation rotates \hat{P} and \hat{Q} around \hat{Z} into \hat{X}, \hat{Y} :

$$\hat{X} = \cos \Omega \hat{P} - \sin \Omega \hat{Q}, \quad (27a)$$

$$\hat{Y} = \sin \Omega \hat{P} + \cos \Omega \hat{Q}, \quad (27b)$$

$$\hat{Z} = \hat{Z}. \quad (27c)$$

The inverse rotation is

$$\hat{P} = \cos \Omega \hat{X} + \sin \Omega \hat{Y}, \quad (28a)$$

$$\hat{Q} = -\sin \Omega \hat{X} + \cos \Omega \hat{Y}, \quad (28b)$$

$$\hat{Z} = \hat{Z}. \quad (28c)$$

Substituting Eqs. (23), (25) and (27) into each other, one obtains the entire transformation:

$$\begin{aligned} \hat{X} &= (\cos \Omega \cos \Phi - \sin \Omega \cos \iota \sin \Phi) \hat{x} \\ &\quad + (-\cos \Omega \sin \Phi - \sin \Omega \cos \iota \cos \Phi) \hat{y} \\ &\quad + \sin \Omega \sin \iota \hat{z}, \end{aligned} \quad (29a)$$

$$\begin{aligned} \hat{Y} &= (\sin \Omega \cos \Phi + \cos \Omega \cos \iota \sin \Phi) \hat{x} \\ &\quad + (-\sin \Omega \sin \Phi + \cos \Omega \cos \iota \cos \Phi) \hat{y} \\ &\quad - \cos \Omega \sin \iota \hat{z}, \end{aligned} \quad (29b)$$

$$\hat{Z} = \sin \iota \sin \Phi \hat{x} + \sin \iota \cos \Phi \hat{y} + \cos \iota \hat{z}. \quad (29c)$$

The inverse transformation is obtained from Eqs. (24), (26) and (28):

$$\begin{aligned} \hat{x} &= (\cos \Omega \cos \Phi - \sin \Omega \cos \iota \sin \Phi) \hat{X} \\ &\quad + (\sin \Omega \cos \Phi + \cos \Omega \cos \iota \sin \Phi) \hat{Y} \\ &\quad + \sin \iota \sin \Phi \hat{Z}, \end{aligned} \quad (30a)$$

$$\begin{aligned} \hat{y} &= (-\cos \Omega \sin \Phi - \sin \Omega \cos \iota \cos \Phi) \hat{X} \\ &\quad + (-\sin \Omega \sin \Phi + \cos \Omega \cos \iota \cos \Phi) \hat{Y} \\ &\quad + \sin \iota \cos \Phi \hat{Z}, \end{aligned} \quad (30b)$$

$$\hat{z} = \sin \Omega \sin \iota \hat{X} - \cos \Omega \sin \iota \hat{Y} + \cos \iota \hat{Z}. \quad (30c)$$

Equation (30c) shows that Ω determines the direction of $\hat{L} = \hat{z}$ on the plane of the sky. For $\Omega = 0$, \hat{L} lies in the \hat{Y} - \hat{Z} plane, and for $\Omega = \pi/2$, it lies in the \hat{X} - \hat{Z} plane. This latter choice ($\Omega = \pi/2$) is already respected by many LAL waveform models. Therefore, in the absence of a reason to do otherwise, all waveform models should default to $\Omega = \pi/2$.

With this recommended default $\Omega = \pi/2$, Eqs. (29) simplify to

$$\hat{X} = -\cos \iota \sin \Phi \hat{x} - \cos \iota \cos \Phi \hat{y} + \sin \iota \hat{z}, \quad (31a)$$

$$\hat{Y} = \cos \Phi \hat{x} - \sin \Phi \hat{y}, \quad (31b)$$

$$\hat{Z} = \sin \iota \sin \Phi \hat{x} + \sin \iota \cos \Phi \hat{y} + \cos \iota \hat{z}. \quad (31c)$$

The inverse transformation (30) simplifies to

$$\hat{x} = -\cos \iota \sin \Phi \hat{X} + \cos \Phi \hat{Y} + \sin \iota \sin \Phi \hat{Z}, \quad (32a)$$

$$\hat{y} = -\cos \iota \cos \Phi \hat{X} - \sin \Phi \hat{Y} + \sin \iota \cos \Phi \hat{Z}, \quad (32b)$$

$$\hat{z} = \sin \iota \hat{X} + \cos \iota \hat{Z}. \quad (32c)$$

2. From NR-frame to wave-frame

Equations (29) and (31) are of particular importance. By definition, the source-frame basis-vectors are trivially related to vectorial quantities of the compact binary dynamics:

$$\hat{x}^{\text{ref}} \equiv \hat{n}, \quad (33a)$$

$$\hat{y}^{\text{ref}} \equiv (\hat{L} \times \hat{n}), \quad (33b)$$

$$\hat{z}^{\text{ref}} \equiv \hat{L}. \quad (33c)$$

Therefore, if the dynamics vectors \hat{L} and \hat{n} are known in any coordinate system, e.g. the NR coordinates, then Eqs. (29) yield the wave-frame basis-vectors in those coordinates. Specifically, for $\Omega = \pi/2$, Eqs. (31) yield

$$\hat{X}^{\text{ref}} \equiv -\cos \iota [\sin \Phi \hat{n} + \cos \Phi \hat{L} \times \hat{n}] + \sin \iota \hat{L}, \quad (34a)$$

$$\hat{Y}^{\text{ref}} \equiv \cos \Phi \hat{n} - \sin \Phi \hat{L} \times \hat{n}, \quad (34b)$$

$$\hat{Z}^{\text{ref}} \equiv \sin \iota [\sin \Phi \hat{n} + \cos \Phi \hat{L} \times \hat{n}] + \cos \iota \hat{L}. \quad (34c)$$

Let us consider next the transformation of the GW strain polarizations from the generic (NR) frame Eq. (11) to the wave-frame. \hat{X} and \hat{Y} are orthogonal to the direction of propagation $\hat{Z} = \hat{r}$ of the gravitational wave. Therefore, $(\hat{\theta}, \hat{\phi})$ can be rotated into (\hat{X}, \hat{Y}) through a rotation by an angle α :

$$\hat{X} = \cos \alpha \hat{\theta} - \sin \alpha \hat{\phi}, \quad (35a)$$

$$\hat{Y} = \sin \alpha \hat{\theta} + \cos \alpha \hat{\phi}. \quad (35b)$$

Substituting Eqs. (35) into Eqs. (22) we can compute h_+^{W} and h_{\times}^{W} in terms of $(\hat{\theta}, \hat{\phi})$. Comparing further with Eqs. (12), we find:

$$h_+^{\text{W}}(t_{\text{GW}}) = \cos(2\alpha) h_+^{\text{NR}}(t_{\text{GW}}) - \sin(2\alpha) h_{\times}^{\text{NR}}(t_{\text{GW}}), \quad (36a)$$

$$h_{\times}^{\text{W}}(t_{\text{GW}}) = \sin(2\alpha) h_+^{\text{NR}}(t_{\text{GW}}) + \cos(2\alpha) h_{\times}^{\text{NR}}(t_{\text{GW}}). \quad (36b)$$

Not surprising, the GW polarizations in the wave-frame are obtained from those in the NR-frame by a rotation of 2α .

The wave-frame \hat{X}, \hat{Y} depend on Ω , and therefore h_+^{W} and h_{\times}^{W} also depend on Ω . We can make this dependence explicit by resorting to the intermediate vectors \hat{P} and \hat{Q} . These are also orthogonal to \hat{Z} , therefore they, too, can be obtained from $\hat{\theta}$ and $\hat{\phi}$ by a rotation:

$$\hat{P} = \cos \alpha' \hat{\theta} - \sin \alpha' \hat{\phi}, \quad (37a)$$

$$\hat{Q} = \sin \alpha' \hat{\theta} + \cos \alpha' \hat{\phi}. \quad (37b)$$

However, \hat{P} and \hat{Q} are independent of Ω and therefore, α' is independent of Ω . Because (\hat{X}, \hat{Y}) are rotated by Ω relative to (\hat{P}, \hat{Q}) , we have

$$\alpha = \Omega + \alpha'. \quad (38)$$

Because rotations add, we can therefore write

$$\begin{pmatrix} h_+^{\text{W}} \\ h_{\times}^{\text{W}} \end{pmatrix} = \mathbf{R}_{2\Omega} \mathbf{R}_{2\alpha'} \begin{pmatrix} h_+^{\text{NR}} \\ h_{\times}^{\text{NR}} \end{pmatrix}, \quad (39)$$

where \mathbf{R}_β denotes a 2x2 rotation matrix,

$$\mathbf{R}_\beta = \begin{pmatrix} \cos \beta & -\sin \beta \\ \sin \beta & \cos \beta \end{pmatrix}. \quad (40)$$

Equation (39) thus implies that the waveform-modes are obtained from the NR-polarizations by (i) applying an Ω -independent rotation by $2\alpha'$; followed by (ii) a rotation by 2Ω . For $\Omega = \pi/2$, we have $\mathbf{R}_\pi = -\mathbf{1}$. Therefore, between $\Omega = 0$ (inclination rotated about X-axis) and $\Omega = \pi/2$ (inclination rotated about Y-axis), the waveform polarization pick up precisely an overall minus-sign.

C. Computing GW polarizations in the LAL wave-frame

Let us finally write down explicit instructions of how to obtain GW polarizations in the LAL convention, given NR waveform data. Given parameters `i`, `phiRef` passed into `XLALSimInspiralChooseTDWaveform`, proceed as follows:

1. Define $\Phi = \text{phiRef}$, $\iota = i$.
2. Compute $\hat{\mathbf{Z}}_{\text{ref}}$ at the reference time by evaluating Eq. (34c).
3. Because $\hat{\mathbf{Z}}$ points in the direction of emission of the gravitational wave, we must have

$$\hat{\mathbf{Z}}_{\text{ref}}^{\text{ref}} \equiv \begin{pmatrix} \cos \phi \sin \theta \\ \sin \phi \sin \theta \\ \cos \theta \end{pmatrix}. \quad (41)$$

From this equality, read off (θ, ϕ) . Then compute the NR-basis vectors $\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}}$ from Eqs. (11).

4. If $\Omega = \pi/2$, compute $\sin \alpha$ and $\cos \alpha$ by taking inner products of Eq. (35b) with $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\phi}}$:

$$\sin \alpha = \cos \Phi \hat{\mathbf{n}} \cdot \hat{\boldsymbol{\theta}} - \sin \Phi (\hat{\mathbf{L}} \times \hat{\mathbf{n}}) \cdot \hat{\boldsymbol{\theta}}, \quad (42a)$$

$$\cos \alpha = \cos \Phi \hat{\mathbf{n}} \cdot \hat{\boldsymbol{\phi}} - \sin \Phi (\hat{\mathbf{L}} \times \hat{\mathbf{n}}) \cdot \hat{\boldsymbol{\phi}}. \quad (42b)$$

If $\Omega \neq \pi/2$, compute instead inner products based on Eqs. (29).

5. Substitute (θ, ϕ) into Eqs. (13) and (12) to compute $h_+^{\text{NR}}(t_{\text{GW}})$ and $h_\times^{\text{NR}}(t_{\text{GW}})$.
6. Compute $\cos 2\alpha = \cos^2 \alpha - \sin^2 \alpha$ and $\sin 2\alpha = 2 \cos \alpha \sin \alpha$. Substitute into Eqs. (36) to compute $h_+^{\text{W}}(t_{\text{GW}})$ and $h_\times^{\text{W}}(t_{\text{GW}})$.

1. Evaluate spin-consistency in LAL source-frame

The parameters `S1x`, `S1y`, `S1z`, `S2x`, `S2y`, `S2z` passed into `LALSimInspiralChooseTDWaveform` are supposed to be the LAL source-frame parameters, i.e. these parameters should simply be the projections of $\vec{\chi}_{1,2}$ onto the source-frame basis-vectors $(\hat{x}, \hat{y}, \hat{z})$. Substituting Eqs. (19) and (20), one arrives at the following consistency conditions:

$$S1x \stackrel{\text{ref}}{=} \vec{\chi}_1 \cdot \hat{\mathbf{n}}, \quad (43a)$$

$$S1y \stackrel{\text{ref}}{=} \vec{\chi}_1 \cdot (\hat{\mathbf{L}} \times \hat{\mathbf{n}}), \quad (43b)$$

$$S1z \stackrel{\text{ref}}{=} \vec{\chi}_1 \cdot \hat{\mathbf{L}}. \quad (43c)$$

The conditions for body 2 are obtained by $1 \leftrightarrow 2$.

V. DISCUSSION

With this new infrastructure it is very easy and much less memory intensive to use NR waveforms directly for data analysis applications. The “NR_hdf5” approximant works much the same as any other approximant in `lalsimulation` but there are a few important differences.

First, the user must supply the location of the HDF5 file, a functionality which was already implemented for NINJA, but was not previously used in `lalsimulation`. Secondly, the user must be careful to supply the mass ratio and spin values that are consistent with the NR files, and the spin values have to be specified in the LAL source frame (see Eqs. (43)).

The current implementation still suffers from a few caveats and drawbacks. As opposed to the continuous waveform approximants, at the moment the metadata are only referring to the beginning of the waveform and not some reference time, which can be chosen freely. While this is not a problem for aligned-spin binaries, this is a big concern for precessing simulations since various quantities, in particular the spins and the orbital angular momentum, are time-dependent. To fully integrate this desired freedom, additional information needs to be incorporated into the HDF5 files *and* the waveform evaluation functions accordingly. Specifically, one needs the time-series of the vectors determining the geometry of the binary: $\hat{L}(t), \hat{n}(t), \chi_1(t), \chi_2(t)$. Given these time-series, one can interpolate these four vectors to any reference epoch, and then apply the frame transformations at this reference epoch. These are provided in the formats 2 and 3 and we leave it to future upgrades to `lalsimulation` to allow for this additional functionality to be fully integrated in the waveform evaluation functions.

Acknowledgements

We are grateful to Mark Hannam for many useful discussions and comments throughout the code review. We also thank Kent Blackburn and James Healy for providing useful comments on the manuscript, and Ian Hinder, Geoffrey Lovelace and Deirdre Shoemaker for input into the metadata discussion. Many thanks for discussions regarding the frame coordinate transformations to Stas Babak, Jolien Creighton, Michael Pürrer and Riccardo Sturani.

-
- [1] B. P. Abbott et al. (Virgo, LIGO Scientific), Phys. Rev. Lett. **116**, 061102 (2016), 1602.03837.
 - [2] B. P. Abbott et al. (Virgo, LIGO Scientific), Phys. Rev. Lett. **116**, 241103 (2016), 1606.04855.
 - [3] B. P. Abbott et al. (Virgo, LIGO Scientific), Phys. Rev. **X6**, 041015 (2016), 1606.04856.
 - [4] B. S. Sathyaprakash and B. F. Schutz, Living Rev. Rel. **12**, 2 (2009), 0903.0338.
 - [5] L. Blanchet, Living Reviews in Relativity **17** (2014), URL <http://www.livingreviews.org/lrr-2014-2>.
 - [6] F. Pretorius, Phys. Rev. Lett. **95**, 121101 (2005), gr-qc/0507014.
 - [7] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter, Phys. Rev. Lett. **96**, 111102 (2006), gr-qc/0511103.
 - [8] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, Phys. Rev. Lett. **96**, 111101 (2006), gr-qc/0511048.
 - [9] J. Centrella, J. G. Baker, B. J. Kelly, and J. R. van Meter, Rev. Mod. Phys. **82**, 3069 (2010), 1010.5260.
 - [10] F. Ohme, Class. Quant. Grav. **29**, 124002 (2012), 1111.3737.
 - [11] M. Hannam, P. Schmidt, A. Boh, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, Phys. Rev. Lett. **113**, 151101 (2014), 1308.3271.
 - [12] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mrou, H. P. Pfeiffer, M. A. Scheel, and B. Szilgyi, Phys. Rev. **D89**, 084006 (2014), 1307.6232.
 - [13] A. Taracchini et al., Phys. Rev. **D89**, 061502 (2014), 1311.2544.

- [14] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Jimnez Forteza, and A. Boh, Phys. Rev. **D93**, 044007 (2016), 1508.07253.
- [15] J. Healy, C. O. Lousto, and Y. Zlochower, Phys. Rev. **D90**, 104004 (2014), 1406.7295.
- [16] B. P. Abbott et al. (Virgo, LIGO Scientific) (2016), 1611.07531.
- [17] B. Aylott et al., Class. Quant. Grav. **26**, 165008 (2009), 0901.4399.
- [18] J. Aasi et al. (VIRGO, LIGO Scientific, NINJA-2), Class. Quant. Grav. **31**, 115004 (2014), 1401.0939.
- [19] C. R. Galley and P. Schmidt (2016), 1611.07529.
- [20] E. Newman and R. Penrose, J. Math. Phys. **3**, 566 (1962).
- [21] R. Penrose, Phys. Rev. Lett. **10**, 66 (1963).
- [22] T. Regge and J. A. Wheeler, Phys. Rev. **108**, 1063 (1957).
- [23] F. J. Zerilli, Phys. Rev. Lett. **24**, 737 (1970).
- [24] F. J. Zerilli, Phys. Rev. **D2**, 2141 (1970).
- [25] V. Moncrief, Annals Phys. **88**, 323 (1974).
- [26] C. R. Galley, *romSpline*, URL <https://bitbucket.org/chadgalley/romspline>.
- [27] D. Brown, S. Fairhurst, B. Krishnan, R. A. Mercer, R. K. Kopparapu, L. Santamaria, and J. T. Whelan (2007), 0709.0093.
- [28] D. R. Fiske, J. G. Baker, J. R. van Meter, D.-I. Choi, and J. M. Centrella, Phys. Rev. **D71**, 104036 (2005), gr-qc/0503100.
- [29] M. Boyle and A. H. Mroue, Phys. Rev. **D80**, 124045 (2009), 0905.3177.
- [30] T. Dal Canton et al., Phys. Rev. **D90**, 082004 (2014), 1405.6731.
- [31] A. H. Mroue et al., Phys. Rev. Lett. **111**, 241104 (2013), 1304.6077.
- [32] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, Phys. Rev. **D49**, 6274 (1994).
- [33] L. E. Kidder, Phys. Rev. **D52**, 821 (1995), gr-qc/9506022.
- [34] URL http://software.ligo.org/docs/lalsuite/lalsimulation/group__lalsimulation__inspiral.html.