ASPECTS OF THE

NUMERICAL SIMULATION

OF BINARY

BLACK HOLE SPACETIMES

AND THE SUITABILITY OF CURRENT APPROACHES TO
THE MODERN ERA OF GRAVITATIONAL WAVE ASTRONOMY

Submitted in Partial Fulfilment of the
Requirements for the Degree of

Doctor of Philosophy

Scientific Research and Thesis Content by

DAVID WILLIAM YEELES

under the Supervision of Professor

MARK D. HANNAM

of the Gravity Exploration Institute

Cardiff University

MMXXI

ABSTRACT

The focus of this thesis is the use of numerical simulations of black hole spacetimes in gravitational wave astronomy. The history of the field over the past century is briefly sketched, tracing the journey from general relativity through to the production of full 3D simulations of the inspiral and merger of a black hole binary. This is followed by an exploration of the current state of the art of numerical black hole binary simulations as used by the LIGO-Virgo-KAGRA Collaboration, first with the presentation of a catalogue of such simulations, then through an investigation into the consistency of simulations produced by three prominent numerical relativity codes. Significant disagreement is observed in their gravitational wave phases which may prove problematic for future ground-based gravitational wave astronomy. Finally the practical obstacles inherent to the simulation of high mass ratio binaries, of particular relevance to upcoming space-based gravitational wave detectors, is discussed and a novel solution proposed with tests on a preliminary implementation of this conducted for static and boosted black holes, and for the head-on collision of two black holes. A reduction in resolution of three orders of magnitude is achieved in these exploratory simulations, equivalent to a speed increase of 60,000% in the static black hole case. While this demonstrates its potential for further development, significant questions remain over how accurately the true solution is recovered or whether it is only the phenomenology that is approximated.

Table of Contents

———— ❖ ————

List of Figures

Preface

—— ◈ ——

The study of gravitational waves dates back to the beginnings of general relativity. That ripples in the spacetime manifold propagate as a wave is a straightforward consequence of the theory, and was accordingly one of the earliest experimental predictions proposed by Einstein in 1916 [1], but for many years opinion shifted back and forth over whether this represented an invariant physical phenomenon or was simply an artifact of the coordinate system used [2]. One thing however was agreed upon: that if physical, these gravitational waves would be far too weak to significantly impact their environment and would therefore remain beyond direct empirical verification. This proved overly pessimistic. The first detection of gravitational waves was announced one hundred years later [3], and as the direct observation of the dynamic geometry of spacetime this event defined the birth of an entirely new branch of astronomy. Throughout history every discovery of a new lens through which to view the universe has resulted in, at the very least, a wealth of scientific discoveries, if not outright paradigm shifts: observations with optical telescopes lead to the overthrow of the geocentric model in the 17$^{\text{th}}$ century, radio astronomy lead to the discovery of cosmic microwave background radiation and hence confirmation of the big bang hypothesis in the 20$^{\text{th}}$ century [4], and most recently in the 21$^{\text{st}}$ century observations of solar neutrinos provided the first tangible evidence of physics beyond the standard model [5]. Gravitational waves serve as an entirely new medium with which to observe, one that is by its nature *clean*, not impeded by dust or stars or galaxies owing to the weak coupling between gravity and matter. In this regard electromagnetic astronomy is fundamentally limited in its scope, only able to look back as far as the moment photons decoupled from matter and the universe became optically transparent approximately 300,000 years after the big bang. The weak coupling to matter opens a window in this surface of last scattering through to the early universe. Similarly,

being purely geometric there is no need to rely on matter to act as the intermediary between gravitational phenomena and observation—black holes can be observed *in vacuo*, providing a direct means to probe gravity in the strong-field environment surrounding black holes and test the extent of general relativity's validity. In this way the observation of gravitational waves can be seen as representing both the final piece of the puzzle to round out the first hundred years of Einsteinian gravitation, and equally the opening of the door to, potentially, hints of some more fundamental underlying mechanism.

Any accelerating mass that lacks axisymmetry will radiate gravitational waves. Even something as prosaic as stirring a cup of tea will produce them, though it took the decaying orbit and merger of two black holes to generate a signal of sufficient strength to be detected in 2015. Such transient signals from the coalescence of two compact objects are currently the only source of gravitational waves we are able to detect, and despite their origin in such extreme environments their influence on Earth is limited to the expansion and contraction of space by roughly one thousandth the diameter of a proton [6]. This therefore requires the construction of exceedingly precise detection apparatus, apparatus that only in recent decades has become technologically and computationally feasible. At the time of writing there are five interferometric gravitational wave detectors in operation: GEO600, LIGO Hanford, LIGO Livingston, Virgo, and KAGRA, of which LIGO and Virgo have in the six years since the first detection published 50 observed gravitational wave events over three observing runs. The fourth such run is expected to add a further 80 to this tally [7] in the next two years, and looking further ahead coming years will see the addition of a sixth detector, LIGO India, to the global network. Plans are already underway for the next generation of gravitational wave observatories, with Einstein Telescope due to come online in the mid-2030s offering unprecedented sensitivity to signals—preliminary estimates suggest on the order of at least 1000 events per year, possibly millions [8]. Together with the proposed LIGO Voyager (2028) and Cosmic Explorer (2035) observatories there will be no shortage of signals, allowing the field of gravitational wave astronomy to mature into a truly statistical science. Working in parallel will be the space-based observatories—DECIGO (2027), Taiji (2033), LISA (2034), TianQin (c. 2030s)—sensitive to much lower-frequency signals and so able to observe black hole binaries for thousands of orbits [9], as well as novel sources such as mergers involving supermassive black holes.

Taken as a whole then gravitational wave astronomy is set to be an increasingly fruitful tool with which to study the universe, with the coming decades offering an abundance of invaluable data. The focus of this thesis spans past, present, and future applications of general relativity to the evolution of dynamic black hole spacetimes for the purposes of gravitational wave astronomy. It is split into three distinct parts that can be considered broadly chronological in this manner.

PART I serves as an introduction to the field of numerical relativity as it has developed from general relativity, and its application to gravitational wave astronomy. The journey from general relativity through the production of black hole simulations to the use of these simulations in the search and analysis of real astrophysical signals is presented, beginning in CHAPTER 1 with an overview of how the Einstein field equations can be formulated in such a way as to permit solution through numerical methods. CHAPTER 2 outlines the representation of black hole spacetimes within this formalism, how they are evolved forwards in time, and how the gravitational radiation content of the spacetime is captured. The practical application of these simulated spacetimes is addressed in CHAPTER 3 along with their place within the wider context of black hole binary signal modelling, and their use in the production of tangible astronomical results. This summary of the field is modelled primarily on the treatment presented in Alcubierre's *Introduction to 3+1 Numerical Relativity* [10].

PART II presents an analysis of the numerical relativity simulations currently in use within gravitational wave astronomy. CHAPTER 4 describes an upcoming catalogue of black hole binary simulations intended for use in the construction of precessing gravitational wave models, spearheaded by Edward Fauchon-Jones but representing the combined effort of the waveforms subgroup at Cardiff's Gravity Exploration Institute. The contributions of the author to this project include the production of eight of the simulations within the catalogue, together with a chapter in the upcoming paper introducing the numerical relativity code used in their production. An analysis of waveform accuracy is presented here for the sake of completeness, adapted from the corresponding chapter written by Eleanor Hamilton for the paper. This chapter is followed by an exploration of the consistency of simulations produced by different numerical relativity codes in CHAPTER 5, a collaborative project that uses simulations produced by the author in addition to Jim Healy of RIT and a team from the SXS collaboration consisting of Katerina Chatzi-

ioannou, Harald Pfeiffer, and Geoffrey Lovelace. Analysis and results are the work of the author.

Finally PART III introduces a novel approach to these numerical relativity simulations that hopes to adapt the field to suit the upcoming era of space-based gravitational wave detectors. It addresses the current inability to produce practical binary simulations in which one component is significantly more massive than the other, a greatly-anticipated source of signals in these future detectors, and proposes a novel solution. The technique is introduced and preliminary tests are performed on the Schwarzschild spacetime in CHAPTER 6, while CHAPTER 7 provides the extension to the boosted case and first attempts at the head-on collision of a binary of mass ratio 1:100. All are found to be stable, though there remains significant work to be done to determine whether these tests have truly recovered an accurate, correct solution to the Einstein field equations or are simply reproducing the correct features. These chapters are solely the work of the author based on an idea by Mark Hannam, with thanks to Jonathan Thompson and Bernard Schutz for valuable discussions.

London, May 2021

Part One

INTRODUCTION
& BACKGROUND

Numerical Relativity

1

With five gravitational wave detectors currently in operation, preparations underway for the next generation, and a brace of space-based observatories on the way gravitational wave astronomy is on a solid footing for the next few decades. But simply building the equipment, as impressive a technical achievement as that is, isn't the end of the story. There is still a significant amount of work necessary to pull out a gravitational wave signal from the data they collect; unlike optical telescopes, where it's simply a matter of looking towards a distant object and watching it shining away in the darkness of space, for these detectors the situation is much more analogous to listening for a single word in a crowded room. In practice the continuous noise-dominated data stream output by these detectors is compared against estimations of what a gravitational wave signal would look like to pick out any signals that may be contained within it and to deduce the properties of the astrophysical system that generated it. These estimates come from models that approximate the morphology of the waveform given some set of source properties, approximations that are themselves derived from the general relativity by solving the field equations representing a spacetime containing two black holes evolving forwards in time. This solution necessarily contains the gravitational wave content of the spacetime.

This first part of the thesis traces back through this journey from general relativity to complete modelled gravitational wave signals ready for use in the field, exploring each of the steps above in detail, where each successive section can be considered an application of the former. The use of general relativity to solve initial value problems is introduced, followed by the application of this to black hole spacetimes, and the implementation of these results in a particular computational framework able to extract the gravitational

radiation. Their role in the construction of waveform models is then discussed, before finally outlining how these models are in turn used for observational astronomy. Our starting point is the essential result at the heart of general relativity: the relation of the geometry of spacetime $G_{\mu\nu}$ to its matter-energy content $T_{\mu\nu}$ through the Einstein field equations[1],

$$G_{\mu\nu} = 8\pi T_{\mu\nu} \ . \qquad\qquad 1\cdot 1$$

The trouble is that this is very hard to solve. It represents a system of ten coupled nonlinear partial differential equations, so an exact solution for anything but the simplest physical scenarios is unattainable. The Schwarzschild solution for example represents the spacetime surrounding an isolated, stationary mass, both spherically symmetric and time-independent. The extension to a rotating mass, the Kerr solution, came only after 48 years of development in the field, and still benefits from simplifications offered by axisymmetry and stationarity. There is little hope for exact solutions for more complex, dynamic systems in astrophysically relevant scenarios with no inherent symmetry.

If the equations cannot be solved analytically then the hope is to solve them numerically. Numerical analysis of differential equations allows approximate solutions to be found even when it is impossible to find an exact one. This is the motivation behind the field of numerical relativity (NR), and amongst its many successes this approach has enabled the evolution of a pair of rotating black holes as they orbit each other, lose energy through the emission of gravitational waves, slowly spiral inwards and eventually merge, leaving a single remnant black hole. While the early stages of the black holes' inspiral can be modelled approximately using Newtonian gravity with additional higher-order correction terms, the later stages necessarily occur in the strong-field, high-velocity regime and as such require a fully relativistic treatment. The inherent complexity of such a system puts it well beyond the reach of analytical general relativity and can only be tackled with the numerical techniques developed over the past 50 years. A great deal of work is necessary to get the Einstein field equations in a form amenable to numerical analysis, and it is with a brief tour of this that we begin.

---

[1]Where throughout Greek is used to denote the spacetime indices of a tensor while Latin indicates spatial, and we use the $(- + ++)$ metric signature.

1.1.   *The 3+1 Formalism*

Early attempts by Hahn in 1963 [11] expanded out the Einstein field equations in some suitable coordinate system and directly evolved the full spacetime metric, representing the binary Schwarzschild spacetime as a single multiply-connected manifold with the horizon of each black hole corresponding to the mouths at either end of a wormhole. This was evolved for 50 time steps "*on an IBM 7090 electronic computer*" using a finite-difference scheme, and showed signs of gravitational attraction between the mouths and collapse of the throat. Many numerical approaches have been taken since then, and among them the most common one, and that used in the simulations discussed in this thesis, is to explicitly introduce an artificial separation between space and time, providing the intuitive interpretation of a system evolving forwards in time. Specifically the aim is to recast the Einstein field equations in the form of an initial value problem, where given the initial state of a system we have a set of equations that can be solved to give its state at some future time.

While there exists a variety of formalisms with which to construct this separation, we will focus here on the *3+1 formalism*, where the spacetime is decomposed into a set of non-interacting, three-dimensional spacelike hypersurfaces, or *slices*[2]. Conceptually this can be pictured as the decomposition of a movie into its stack of component frames, each one capturing a single moment that, considered together, form an evolving picture. Indeed such a decomposition is called a *foliation* of the spacetime, formally defined for a given manifold M as a set of three-dimensional submanifolds Σ_t such that some smooth scalar field t defined on M has non-zero gradient everywhere. Any globally hyperbolic spacetime—that is, one with a Cauchy surface, an achronal set of points causally connected to the entire spacetime—can be foliated in this fashion, sufficient for the scenarios under consideration in this thesis.

Our aim is to reformulate the Einstein field equations in this 3 + 1 decomposition. The key feature of these equations is the curvature of the spacetime manifold, encoded by the four-dimensional Riemann tensor $R_{\alpha\beta\mu\nu}$. Following our foliation of the spacetime we must now define both the *intrinsic* curvature of the spatial slice, $\overset{3}{R}_{\alpha\beta\mu\nu}$, and in addition the *extrinsic* curvature that arises from the way the slices are embedded in the bulk, $K_{\mu\nu}$. Intrinsic curvature is simply the on-slice Riemann tensor, but extrinsic curvature

---

[2]cf. the 1+3 'threading' formalism in which it is foliated by timelike curves [12]

is determined by parallel transporting the normal to the slices n^μ across that slice, or equivalently taking the Lie derivative of the on-slice metric γ_{ij} along the normal to the slice (see [10]),

$$K_{\mu\nu} = -\frac{1}{2}\mathcal{L}_{\vec{n}}\gamma_{\mu\nu} \, ,$$

<div align="right">1·2</div>

and can therefore be conceptualised as the change in the spatial metric as normal observers move forwards in time. The decomposition of the full Riemann tensor into its projection on to a spatial slice and the extrinsic curvature is given by the Gauss-Codazzi equations [10],

$$\mathcal{P}_\alpha^\delta \mathcal{P}_\beta^\kappa \mathcal{P}_\mu^\lambda \mathcal{P}_\nu^\sigma R_{\delta\kappa\lambda\sigma} = \overset{3}{R}_{\alpha\beta\mu\nu} + K_{\alpha\mu}K_{\beta\nu} - K_{\alpha\nu}K_{\beta\mu} \, ,$$

<div align="right">1·3</div>

where the slice-projection operator is defined as $\mathcal{P}^{\mu\nu} \equiv (g^\mu_\nu + n^\mu n_\nu)$. We can contract this to give

$$\mathcal{P}^{\alpha\mu}\mathcal{P}^{\beta\nu}R_{\alpha\beta\mu\nu} = \overset{3}{R} + K^2 - K_{\mu\nu}K_{\mu\nu} \, ,$$

<div align="right">1·4</div>

but if we simply compute the left-hand side from scratch we find that

$$\mathcal{P}^{\alpha\mu}\mathcal{P}^{\beta\nu}R_{\alpha\beta\mu\nu} = R + 2n^\mu n^\nu R_{\mu\nu} \, .$$

<div align="right">1·5</div>

Together then we have that

$$\begin{aligned} \overset{3}{R} + K^2 - K_{\mu\nu}K_{\mu\nu} &= R + 2n^\mu n^\nu R_{\mu\nu} \, , \\ &= 2n^\mu n^\nu G_{\mu\nu} \, , \\ &= 16\pi n^\mu n^\nu T_{\mu\nu} \, , \\ &= 16\pi\rho \, , \end{aligned}$$

<div align="right">1·6</div>

where we have used the standard Einstein tensor $G_{\mu\nu}$ and defined ρ as the energy density as measured by normal observers. Note the lack of any time dependence in 1·6: this equation tells us nothing about how the system evolves but is rather a condition that must at all times be satisfied. It is known as the *Hamiltonian constraint*.

Now if we take the projection of the curvature on the slice again, but this time contract it with the normal vector, we arrive at the Codazzi-Mainardi equations [10]

$$
\begin{aligned}
\mathcal{P}_\alpha^\delta \mathcal{P}_\beta^\kappa \mathcal{P}_\mu^\lambda n^\nu R_{\delta\kappa\lambda\nu} &= \mathcal{P}_\beta^\varepsilon \nabla_\varepsilon K_{\alpha\mu} - \mathcal{P}^\tau{}_\alpha \nabla_\tau K_{\beta\mu} \,, \\
&= D_\beta K_{\alpha\mu} - D_\alpha K_{\beta\mu} \,,
\end{aligned}
$$

<div align="right">1·7</div>

where ∇ is the full four-dimensional spacetime covariant derivative while D is the three-dimensional version on the slice, $D_\beta \equiv \mathcal{P}_\beta^\varepsilon \nabla_\varepsilon$. Under contraction this gives

$$
\begin{aligned}
\mathcal{P}^{\alpha\mu} n^\nu R_{\mu\nu} &= D^\alpha K - D_\mu K^{\alpha\mu} \,, \\
&= D_\mu \left(\gamma^{\alpha\mu} K - K^{\alpha\mu} \right) \,.
\end{aligned}
$$

<div align="right">1·8</div>

We can then substitute $R_{\mu\nu}$ using the Einstein field equations once again to find

$$
\begin{aligned}
D_\mu \left(K^{\alpha\mu} - \gamma^{\alpha\mu} \right) &= -\mathcal{P}^{\alpha\mu} n^\nu \left(G_{\mu\nu} + \frac{1}{2} R g_{\mu\nu} \right) \,, \\
&= -\mathcal{P}^{\alpha\mu} n^\nu G_{\mu\nu} \,, \\
&= -8\pi \mathcal{P}^{\alpha\mu} n^\nu T_{\mu\nu} \,, \\
&\equiv 8\pi j^\alpha \,.
\end{aligned}
$$

<div align="right">1·9</div>

where j^α is the momentum density as measured by normal observers. This is the *momentum constraint*. If we restrict our focus to the spatial parts of these constraint equations we find that on each slice we must satisfy

$$
\begin{aligned}
\overset{3}{R} + K^2 - K_{ij} K_{ij} &= 16\pi\rho \,, \\
D_j \left(K^{ij} - \gamma^{ij} K \right) &= 8\pi j^i \,.
\end{aligned}
$$

<div align="right">1·10</div>

But these represent just four of the ten Einstein field equations. The rest encode the evolution of the gravitational field, and are contained within two evolution equations we are yet to obtain: one for the spatial metric, and one for the extrinsic curvature.

To obtain the first of these we take a step back and consider two adjacent slices Σ_t and Σ_{t+dt} covered by the coordinate lines x^i, as shown in , and with coordinate time

$$
t^\mu = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} \,.
$$

<div align="right">1·11</div>

For an observer moving along the normal to the slices the lapse in coordinate time is trivially $t^\mu dt$, while the lapse in proper time between Σ_t and Σ_{t+dt} is given by

$$d\tau = \alpha(t, x^i)dt \; , \qquad\qquad\qquad 1\cdot12$$

and their shift in position relative to the coordinate grid is given by

$$dx^i = -\beta^i(t, x^i)dt \; . \qquad\qquad\qquad 1\cdot13$$

These two functions α and β^i are accordingly known as the *lapse* and the *shift vector*, and the four-dimensional line element written in terms of them is

$$ds^2 = -\alpha^2 dt^2 + \gamma_{ij}(dx^i + \beta^i dt)(dx^j + \beta^j dt) \; . \qquad\qquad\qquad 1\cdot14$$

From FIG 1.1 we have then through simple vector addition that

$$t^\mu dt = \alpha n^\mu dt + \beta^i dt \; , \qquad\qquad\qquad 1\cdot15$$

which can be rearranged to give

$$n^\mu = \frac{t^\mu}{\alpha} - \frac{\beta^i}{\alpha} \; . \qquad\qquad\qquad 1\cdot16$$

This, together with 1·11, gives us the components of the normal to the slices

$$n^\mu = \begin{pmatrix} n^0 \\ n^i \end{pmatrix} = \begin{pmatrix} \frac{1}{\alpha} \\ -\frac{\beta^i}{\alpha} \end{pmatrix} \; . \qquad\qquad\qquad 1\cdot17$$

The key here is that the specific way in which you foliate your spacetime is entirely unconstrained: α and β^i are freely chosen gauge functions, encoding the inherent gauge freedom of general relativity.

With all this introduced, if we now return to 1·2 notice that we are free to change the vector along which we take the Lie derivative: for some scalar s, the Lie derivative of some tensor $X_{\mu\nu}$ with respect to the vector $\vec{V}$ is

$$\mathcal{L}_{s\vec{V}} X_{\mu\nu} = s\mathcal{L}_{\vec{V}} X_{\mu\nu} \; . \qquad\qquad\qquad 1\cdot18$$
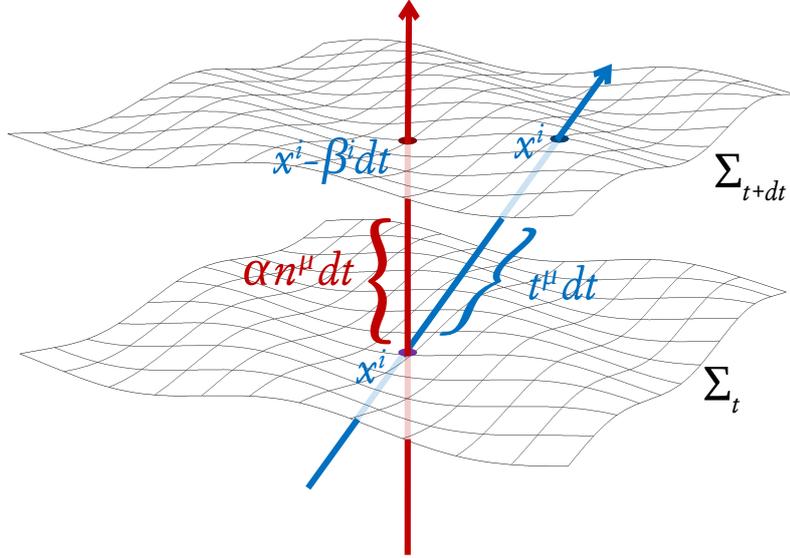
FIGURE 1.1   *Two consecutive three-dimensional slices. The red line is normal to the slices while the blue is a line of constant spatial coordinate. The time step between the slices is quantified by the lapse function α while the spatial coordinate drift is given by the shift vector β^i.*

Taking the lapse as just such a scalar, 1·2 becomes

$$K_{\mu\nu} = -\frac{1}{2\alpha}\mathcal{L}_{\alpha\vec{n}}\gamma_{\mu\nu}\,, \qquad\qquad 1\cdot19$$

and given 1·16 we can substitute $\mathcal{L}_{\alpha\vec{n}} = \mathcal{L}_{\vec{t}} - \mathcal{L}_{\vec{\beta}}$ to rewrite this as

$$\mathcal{L}_{\vec{t}}\gamma_{\mu\nu} = -2\alpha K_{\mu\nu} + \mathcal{L}_{\vec{\beta}}\gamma_{\mu\nu}\,. \qquad\qquad 1\cdot20$$

The particular coordinate system we are employing here, specifically 1·11, means we have simply $\mathcal{L}_t = \partial_t$. As before, taking only the spatial part of this to focus on the behaviour on the slice and correspondingly replacing the Lie derivatives with the on-slice spatial covariant derivatives D^i, we arrive at

$$\partial_t\gamma_{ij} = -2\alpha K_{ij} + D_i\beta_j + D_j\beta_i\,, \qquad\qquad 1\cdot21$$

an evolution equation for the spatial metric. All we need now to complete our reformulation of the Einstein equations as an initial value problem is the corresponding one for the extrinsic curvature. We previously explored $\mathcal{P}_\alpha^\delta \mathcal{P}_\beta^\kappa \mathcal{P}_\mu^\lambda \mathcal{P}_\nu^\sigma R_{\delta\kappa\lambda\sigma}$ and $\mathcal{P}_\alpha^\delta \mathcal{P}_\beta^\kappa \mathcal{P}_\mu^\lambda n^\nu R_{\delta\kappa\lambda\nu}$ (1·3 and 1·7), and to obtain this we consider now $\mathcal{P}_\mu^\delta \mathcal{P}_\nu^\kappa n^\lambda n^\sigma R_{\delta\lambda\kappa\sigma}$, which turns out to be [13]

$$
\begin{aligned}
\mathcal{P}_\mu^\delta \mathcal{P}_\nu^\kappa n^\lambda n^\sigma R_{\delta\lambda\kappa\sigma} &= \mathcal{L}_{\vec{n}} K_{\mu\nu} + K_{\mu\lambda} K_\nu^\lambda + \frac{1}{\alpha} D_\mu D_\nu \alpha \,, \\
&= \left(\mathcal{L}_t - \mathcal{L}_{\vec{\beta}} \right) K_{\mu\lambda} + K_{\mu\lambda} K_\nu^\lambda + \frac{1}{\alpha} D_\mu D_\nu \alpha \,.
\end{aligned}
\tag{1·22}
$$

We can also expand a couple of the projection operators in 1·3 to give

$$
\mathcal{P}_\mu^\delta \mathcal{P}_\nu^\kappa n^\lambda n^\sigma \left(R_{\delta\lambda\kappa\sigma} + R_{\delta\kappa} \right) = \overset{3}{R}_{\mu\nu} + K K_{\mu\nu} - K_{\mu\lambda} K_\nu^\lambda \,,
\tag{1·23}
$$

and together we can use these to construct

$$
\begin{aligned}
\left(\mathcal{L}_t - \mathcal{L}_{\vec{\beta}} \right) K_{\mu\nu} &= -D_\mu D_\nu \alpha + \alpha \left(\overset{3}{R}_{\mu\nu} + K K_{\mu\nu} - 2 K_{\mu\lambda} K_\nu^\lambda \right) + \alpha \left(-\mathcal{P}_\mu^\delta \mathcal{P}_\nu^\kappa R_{\delta\kappa} \right) \,, \\
&= \quad\quad '' \quad\quad\quad\quad + \alpha \left(-\mathcal{P}_\mu^\delta \mathcal{P}_\nu^\kappa \left(8\pi T_{\mu\nu} - \frac{1}{2} \gamma_{\mu\nu} T \right) \right) \,, \\
&= \quad\quad '' \quad\quad\quad\quad + 4\pi\alpha \left(\mathcal{P}_\mu^\delta \mathcal{P}_\nu^\kappa \gamma_{\mu\nu} T - 2 \mathcal{P}_\mu^\delta \mathcal{P}_\nu^\kappa T_{\mu\nu} \right) \,, \\
&= \quad\quad '' \quad\quad\quad\quad + 4\pi\alpha \left(\gamma_{\mu\nu} \left(S - \rho \right) - 2 S_{\mu\nu} \right) \,,
\end{aligned}
\tag{1·24}
$$

where we have used the definition of the energy density given in 1·6 and introduced $S_{\mu\nu} \equiv \mathcal{P}_\mu^\alpha \mathcal{P}_\nu^\beta T_{\alpha\beta}$, the spatial stress as measured by normal observers. The on-slice contribution to this is

$$
\begin{aligned}
\partial_t K_{ij} &= - \mathcal{L}_{\vec{\beta}} K_{ij} - D_i D_j \alpha \\
&\quad + \alpha \left(\overset{3}{R}_{ij} + K K_{ij} - 2 K_{ik} K_j^k \right) + 4\pi\alpha \left(\gamma_{ij}(S - \rho) - 2 S_{ij} \right) \,, \\
&= \beta^k \partial_k K_{ij} + K_{ki} \partial_j \beta^k + K_{kj} \partial_i \beta^k - D_i D_j \alpha \\
&\quad + \alpha \left(\overset{3}{R}_{ij} + K K_{ij} - 2 K_{ik} K_j^k \right) + 4\pi\alpha \left(\gamma_{ij}(S - \rho) - 2 S_{ij} \right) \,,
\end{aligned}
\tag{1·25}
$$

Together 1·21 and 1·25 make up the Arnowitt-Deser-Misner (ADM) equations. These are by no means unique—after all, you can freely add multiples of the constraints to them without affecting the solution (because by definition they will vanish for any valid solution to the Einstein equations) and so build up different evolution equations. A similar freedom exists in the choice of conformal scaling and tensor splittings, and the introduction of auxiliary variables. One such alternative that takes advantage of this freedom and has proved extremely successful in the numerical relativity community is that first presented by Nakamura, Oohara and Kojima [14], refined by Shibata and Nakamura [15] and popularised by Baumgarte and Shapiro [16], known as the BSSN formulation. This enjoys a long-term stability that ADM lacks, and as a result has become a common choice for the evolution of black hole spacetimes.

1.2. *The BSSN Formalism*

The derivation presented here follows that given in chapters 2 and 3 of [10]. The BSSN formulation is based primarily upon a conformal rescaling of the ADM equations. We take the conformal transformation of the on-slice metric

$$\tilde{\gamma}_{ij} = \psi^{-4}\gamma_{ij} \, , \tag{1·26}$$

and further introduce the conformal connection on the slice defined by

$$\widetilde{\Gamma}^i = \tilde{\gamma}^{jk}\widetilde{\Gamma}^i_{jk} \, . \tag{1·27}$$

The only requirement placed on the conformal factor ψ is that it be chosen such that $\tilde{\gamma}$ = det $\tilde{\gamma}_{ij}$ = 1, which means 1·27 becomes

$$\widetilde{\Gamma}^i = -\partial_j\tilde{\gamma}^{ij} \, . \tag{1·28}$$

A further distinction of the BSSN formulation from the ADM is the decomposition of the extrinsic curvature into its trace K and its traceless part A_{ij},

$$K_{ij} = A_{ij} + \frac{1}{3}\gamma_{ij}K \, , \tag{1·29}$$

which itself is conformally transformed into

$$\widetilde{A}_{ij} = \psi^{-4} A_{ij} \, . \qquad\qquad 1 \cdot 30$$

With all this reordering and introduction of new variables, rather than just the two evolution equations for K_{ij} and γ_{ij} we now need evolution equations for K, $\widetilde{A}_{ij}$, $\widetilde{\gamma}_{ij}$, $\widetilde{\Gamma}^i$, and ψ. The first of these come directly from the ADM equations 1·21 and 1·25:

$$\partial_t K = -D_i D^i \alpha + \alpha \left(\widetilde{A}_{ij} \widetilde{A}^{ij} + \frac{1}{3} K^2 \right) + 4\pi\alpha(\rho + S) \, , \qquad\qquad 1 \cdot 31$$

$$\partial_t \widetilde{A}_{ij} = \psi^{-4} \underbrace{\left(-D_i D_j \alpha + \alpha R_{ij} + 4\pi\alpha \left(\gamma_{ij}(S - \rho) - 2 S_{ij} \right) \right)}_{traceless} + \alpha \left(K \widetilde{A}_{ij} - 2 \widetilde{A}_{ik} \widetilde{A}^k_j \right) \, , \quad 1 \cdot 32$$

$$\partial_t \widetilde{\gamma}_{ij} = -2\alpha \widetilde{A}_{ij} \, , \qquad\qquad 1 \cdot 33$$

and from 1·21 we can find the evolution equation for $\gamma \equiv \det \gamma_{ij}$

$$\begin{aligned}
\partial_t \gamma &= \gamma \left(-2\alpha K + 2 D_i \beta^j \right) \, , \\
&= -2\gamma\alpha K + 2\gamma \partial_i \beta^i + \beta^i \partial_i \gamma \, ,
\end{aligned} \qquad\qquad 1 \cdot 34$$

which given the requirement that $\widetilde{\gamma} = 1$ (i.e. $\psi = \gamma^{-12}$) gives us that

$$\begin{aligned}
\partial_t \psi &= -\frac{\psi}{6} \left(\alpha K - \partial_i \beta^i \right) + \beta^i \partial_i \psi \, , \\
&= -\frac{1}{6} \alpha K \psi \, .
\end{aligned} \qquad\qquad 1 \cdot 35$$

That just leaves us yet to find the evolution equation for $\widetilde{\Gamma}^i$. From its definition in 1·28, together with the γ_{ij} evolution equation in 1·20, we have that

$$\begin{aligned}
\partial_t \widetilde{\Gamma}^i &= \partial_t \left(-\partial_j \widetilde{\gamma}^{ij} \right) \, , \\
&= -\partial_j \left(\partial_t \widetilde{\gamma}^{ij} \right) \, , \\
&= -\partial_j \left(\mathcal{L}_t \widetilde{\gamma}^{ij} \right) \, , \\
&= -\partial_j \left(\mathcal{L}_{\vec{\beta}} \widetilde{\gamma}^{ij} + 2\alpha K^{ij} \right) \, ,
\end{aligned}$$

$$
\begin{aligned}
&= -\partial_j \left(\mathcal{L}_{\tilde{\beta}} \tilde{\gamma}^{ij} + 2\alpha \widetilde{A}^{ij} \right) , \\
&= -\partial_j \left(\mathcal{L}_{\tilde{\beta}} \tilde{\gamma}^{ij} \right) - 2 \left(\alpha \partial_j \widetilde{A}^{ij} + \widetilde{A}^{ij} \partial_j \alpha \right) , \\
&= \tilde{\gamma}^{jk} \partial_j \partial_k \beta^i + \frac{1}{3} \tilde{\gamma}^{ij} \partial_j \partial_k \beta^k - 2 \left(\alpha \partial_j \widetilde{A}^{ij} + \widetilde{A}^{ij} \partial_j \alpha \right) ,
\end{aligned}
$$

<div align="right">1·36</div>

but unfortunately this lacks stability due to the divergence of the traceless extrinsic curvature. The solution lies in replacing this using the momentum constraint 1·10, which in terms of our BSSN variables is

$$
\partial_j \widetilde{A}^{ij} = -\widetilde{\Gamma}^i_{jk} \widetilde{A}^{jk} - \frac{6}{\psi} \widetilde{A}^{ij} \partial_j \psi + \frac{2}{3} \tilde{\gamma}^{ij} \partial_j K + 8\pi \tilde{j}^i ,
$$

<div align="right">1·37</div>

leaving us with

$$
\begin{aligned}
\partial_t \widetilde{\Gamma}^i = {}& \tilde{\gamma}^{jk} \partial_j \partial_k \beta^i + \frac{1}{3} \tilde{\gamma}^{ij} \partial_j \partial_k \beta^k - 2\tilde{\gamma}^{ij} \partial_j \alpha \\
& + 2\alpha \left(\widetilde{\Gamma}^i_{jk} \widetilde{A}^{jk} + \frac{6}{\psi} \widetilde{A}^{ij} \partial_j \psi - \frac{2}{3} \tilde{\gamma}^{ij} \partial_j K - 8\pi \tilde{j}^i \right) ,
\end{aligned}
$$

<div align="right">1·38</div>

which completes the set of five BSSN evolution equations.

With these in hand together with our constraint equations we now have everything needed to start evolving a system. All that remains is to provide some initial state and then solve the evolution equations for each future time. The particular way in which this initial state is chosen is the focus of the next section, where we see how to solve the constraint equations for a variety of black hole spacetimes.

1.3. *Initial Data Construction*

Having discovered two constraints that must be satisfied at all times by our spacetime it is clear that we can't specify the initial geometry arbitrarily, and must instead solve the constraint equations to provide the appropriate γ_{ij}, K_{ij} for a given physical scenario. This is certainly not trivial given they form a system of four coupled elliptic partial differential equations, but a number of techniques [17–19] have been developed to tackle this problem. We focus here on the *York-Lichnerowicz conformal (transverse) decomposition* [20, 21].

We begin by taking the Hamiltonian and momentum constraints 1·10 and casting them in BSSN form,

$$8\widetilde{D}^2\psi - \overset{3}{\widetilde{R}}\psi + \psi^5\left(A_{ij}A^{ij} - \frac{2}{3}K^2\right) + 16\pi\psi^5\rho = 0 \,, \qquad 1\cdot39$$

$$D_j A^{ij} - \frac{2}{3}D^i K - 8\pi j^i = 0 \,, \qquad 1\cdot40$$

which has transformed the Hamiltonian constraint into an elliptic equation in ψ. If we solve this for ψ we will be able to reconstruct the physical on-slice metric γ_{ij} from some conformal $\widetilde{\gamma}_{ij}$. As for the momentum constraint, we have the three unknowns A^{ij}, K, j^i and so require three equations that can be solved to obtain them.

At this point we change slightly our conformal transformation in order to simplify the final expressions we derive. The covariant derivative of a symmetric traceless tensor, for example our traceless extrinsic curvature A^{ij}, obeys [10]

$$D_j A^{ij} = \frac{1}{\psi^n}D_j(\psi^n A^{ij}) + (10-n)A^{ik}\partial_k\ln\psi \,, \qquad 1\cdot41$$

so by using the conformal factor ψ^{10} rather than ψ^4 as we have used up to this point we can eliminate the second term. A further property of symmetric traceless tensors is that they can be decomposed into their transverse and longitudinal parts:

$$A^{ij} = \underbrace{A_*^{ij}}_{transverse} + \underbrace{(\mathscr{L}W)^{ij}}_{longitudinal} \,, \qquad 1\cdot42$$

with the longitudinal part the conformal Killing form of the vector $\vec{W}$, generated by the operator

$$(\mathscr{L}W)^{ij} \equiv D^i W^j + D^j W^i - \frac{2}{3}\gamma^{ij}D_k W^k \,. \qquad 1\cdot43$$

This gives the momentum constraint in the form of a set of three coupled equation for the vector W^i:

$$\widetilde{D}_j(\mathscr{L}W)^{ij} - \frac{2}{3}\psi^6\widetilde{D}^i K - 8\pi\psi^{10}j^i = 0 \,. \qquad 1\cdot44$$

With our refashioned constraint equations we can now compute ψ and W^i (and hence γ_{ij} and K^{ij}) given some $\widetilde{\gamma}_{ij}$, K, and $\widetilde{A}_*^{ij}$.

The trouble is that while it's straightforward to construct a symmetric traceless tensor like $\widetilde{A}^{ij}$ it's much more difficult to ensure it is transverse, satisfying $D_i \widetilde{A}^{ij}_*$. To tackle this we introduce two auxiliary quantities, some symmetric traceless tensor $\widetilde{M^{ij}}$ and a vector $\widetilde{Y}^i$. As before, $\widetilde{M}^{ij}$ can be decomposed as

$$\widetilde{M}^{ij} = \widetilde{M}^{ij}_* - (\mathscr{L}Y)^{ij} \, , \tag{1.45}$$

and as $\widetilde{M}^{ij}_*$ is by definition transverse it must follow that

$$\widetilde{D}_j \, (\mathscr{L}Y)^{ij} = \widetilde{D}_j \widetilde{M}^{ij} \, , \tag{1.46}$$

which, given $\widetilde{M^{ij}}$, can be solved for $\widetilde{Y}^i$ and in turn give us $\widetilde{M}^{ij}_*$.

If we then take this to be the transverse extrinsic curvature that we're trying to find (i.e. $\widetilde{M}^{ij}_* = \widetilde{A}^{ij}_*$), we then have from 1.42

$$
\begin{aligned}
\widetilde{A}^{ij} &= \widetilde{A}^{ij}_* + (\mathscr{L}W)^{ij} \, , \\
&= \widetilde{M}^{ij}_* + (\mathscr{L}W)^{ij} \, , \\
&= \left(\widetilde{M}^{ij} - (\mathscr{L}Y)^{ij} \right) + (\mathscr{L}W)^{ij} \, , \\
&= \widetilde{M}^{ij} + \left(- (\mathscr{L}Y)^{ij} + (\mathscr{L}W)^{ij} \right) \, , \\
&= \widetilde{M}^{ij} + (\mathscr{L}V)^{ij} \, ,
\end{aligned} \tag{1.47}
$$

where $\widetilde{V}^i \equiv \widetilde{W}^i - \widetilde{Y}^i$.

Finally then, rather than being given initial data in the form of $\widetilde{A}^{ij}_*$ we are instead given $\widetilde{A}^{ij}$, $\widetilde{M}^{ij}$, and $\widetilde{V}^i$. Correspondingly the momentum constraint is written

$$\widetilde{D}_j \, (\mathscr{L}V)^{ij} + \widetilde{D}_j \widetilde{M}^{ij} - \frac{2}{3} \psi^6 \widetilde{D}^i K - 8\pi \psi^{10} j^i = 0 \, . \tag{1.48}$$

The constraint equations are then solved for ψ and $\widetilde{V}^i$, with free data specified in $\widetilde{\gamma}_{ij}$, $\widetilde{M}^{ij}$, K, $\widetilde{\rho}$, and $\widetilde{j}^i$, and finally the physical quantities can be reconstructed as

$$\gamma_{ij} = \psi^4 \widetilde{\gamma}_{ij} \, , \tag{1.49}$$

$$K^{ij} = \psi^{-10} \left(\widetilde{M}^{ij} + \widetilde{D}_j \, (\mathscr{L}V)^{ij} \right) + \frac{1}{3} \gamma^{ij} K \, . \tag{1.50}$$

These variables have no obvious interpretation, but if one adopts the *conformal thin-sandwich* decomposition in which the conformal metric is specified on two neighbouring slices then they can be related to the shift and conformal metric as [10]

$$\widetilde{V}^i = \beta^i \, , \tag{1.51}$$

$$\widetilde{M}^{ij} = -\frac{1}{2\alpha} \partial_t \widetilde{\gamma}^{ij} \, . \tag{1.52}$$

At this stage we have all that we need to begin evolving a given spacetime. We have a rigorous way to define a physically valid initial state of the spacetime, along with a set of partial differential equations that, when solved, return its future state. All of this is phrased in terms of quantities that have a relatively straightforward physical interpretation, dependent only on the slice geometry (K and γ) and the coordinate structure (α, β, ψ). In the next chapter we apply these general results to the specific case of black hole spacetimes, translating some key results of general relativity into the framework we have introduced here.

Simulating Black Hole Spacetimes

2

So far we have explored the recasting of general relativity in a form suitable for numerical analysis in a general sense, discussing the methods used and deriving key results in a fairly scenario-agnostic manner aside from the requirement of global hyperbolicity for the foliation of our spacetime. We turn now to the primary focus of this thesis and indeed the primary application of numerical relativity in the context of gravitational wave astronomy: black hole binaries. This in fact allows us to simplify the picture significantly, particularly in the illustrative, if somewhat contrived, example of two momentarily-static Schwarzschild black holes. We begin with just one of these and end up at multiple Kerr black hole initial data, before moving on to a discussion of some of the more practical aspects involved in their evolution such as the extraction of the dynamics. We conclude with a brief introduction to one particular implementation of all that has been discussed in these first two chapters: the BAM code.

2.1. *Schwarzschild Initial Data*

For a static ($K_{ij} = 0$) conformally flat ($\overset{3}{\tilde{R}} = 0$) vacuum ($\rho = 0$) spacetime upon which we impose that $K = 0$ (the *maximal slicing* condition), the Hamiltonian constraint 1·39 reduces to simply

$$D^2 \psi = 0 \ . \tag{2·1}$$

Imposing the natural boundary condition that the spacetime is asymptotically flat ($\psi|_\infty = 1$), this is trivially solved by simply $\psi = 1$, which would return the on-slice metric $\gamma_{ij} =$

$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$: flat space. The simplest *nontrivial* solution would be, for the radial coordinate r,

$$\psi = 1 + \frac{k}{r} \,, \qquad\qquad\qquad 2\cdot2$$

giving the spatial metric

$$\gamma_{ij} = \begin{pmatrix} \left(1 + \frac{k}{r}\right)^4 & 0 & 0 \\ 0 & \left(1 + \frac{k}{r}\right)^4 r^2 & 0 \\ 0 & 0 & \left(1 + \frac{k}{r}\right)^4 r^2 \sin^2\theta \end{pmatrix} \,, \qquad 2\cdot3$$

which is the spatial part of the Schwarzschild metric in isotropic coordinates with $\frac{M}{2} \to k$. The temporal part of the metric in these coordinates is given by the lapse

$$\alpha = \frac{1 - \frac{M}{2r}}{1 + \frac{M}{2r}} \,. \qquad\qquad\qquad 2\cdot4$$

Thus we have complete initial data for a single Schwarzschild black hole. But of course in 2·1 we can take advantage of the linearity of Laplace's equation and form new solutions by adding together other solutions. For two black holes then we can simply take

$$\begin{aligned} \psi &= 1 + \frac{k_1}{|r - r_1|} + \frac{k_2}{|r - r_2|} \,, \\ &= 1 + \frac{M_1}{2|r - r_1|} + \frac{M_2}{2|r - r_2|} \,. \end{aligned} \qquad 2\cdot5$$

The general form of this for N black holes is known as the Brill-Lindquist initial data [22, 23]. Note that the singular points r_i represent spatial infinity in *different* asymptotically flat regions, so the spacetime considered as a whole is a multiply-connected manifold consisting of N separate universes joined by wormholes to the 'prime' universe. The points r_i themselves are however not part of the manifold, as can be seen from the sketch in FIG 2.1, so we consider the solution to be given in $\mathbb{R}^3$ *punctured* by N points. This rather exotic topology may seem somewhat troubling at first, but it is all contained within the horizon and so cannot influence the rest of the spacetime.
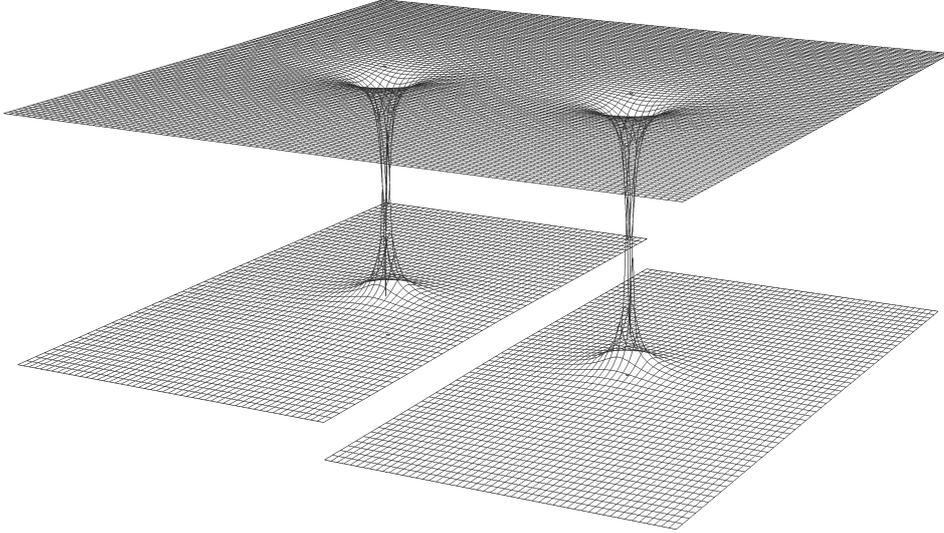
FIGURE 2.1 *Schematic illustrating the nontrivial topology of the Brill-Lindquist initial data. The two black holes in the upper region are represented by wormholes connecting to the two lower additional asymptotically flat regions. In this way the horizon corresponds to the throat of the wormhole while the singularity corresponds to spatial infinity in the secondary universe.*

2.2. *Kerr Initial Data*

The preceding discussion of the Schwarzschild black hole benefitted from the drastic simplifications that come with temporal symmetry, but it is rather unphysical—astrophysical black holes will both be moving and are generically expected to be spinning, so we need to construct initial data that are asymmetric in time and include these features. We begin with some of the same simplifications we used before, namely that we work with a maximally-sliced conformally flat vacuum. Our free parameter $\widetilde{M}^{ij}$ we set to zero, and so the momentum constraint 1·48 reduces to

$$\widetilde{D}_j \left(\mathscr{L} V \right)^{ij} = 0 \,.$$

2·6

One solution[1] to this that produces an extrinsic curvature for a boosted Kerr black hole is given by [10]

$$\widetilde{V}^i = -\frac{1}{4r} \left(7 P^i + n^i n_j P^j \right) + \frac{1}{r^2} \varepsilon^{ijk} n_j S_k \,,$$

2·7

---

[1]See [24] for a more complete list of solutions.

for some constant vectors P^i and S^i, and where we now use n^i to signify the outward unit radial vector rather than the unit normal to the slice. We now have our $\widetilde{V}^i$ and so, following 1·47, can construct $\widetilde{A}^{ij}$ as

$$\widetilde{A}^{ij} = \frac{3}{2r^2}\left(n_i P_j + n_j P_i + n_k P^k\left(n_i n_j - \delta_{ij}\right)\right) - \frac{3}{r^3}\left(\varepsilon_{ijk}n_j + \varepsilon_{jlk}n_i\right)n^l S^k\ , \qquad 2\cdot 8$$

and through 1·49 can reconstruct the physical extrinsic curvature by simply multiplying by ψ^{-2}. This is the *Bowen-York extrinsic curvature*. If we substitute this into the ADM momenta integrals derived in [13],

$$P^i_{ADM} = \frac{1}{8\pi}\lim_{r\to\infty}\oint\left(K^i_l - \delta^i_l K\right)n^l\,dS\ , \qquad 2\cdot 9$$

$$J^i_{ADM} = \frac{1}{16\pi}\lim_{r\to\infty}\oint\varepsilon^{ijk}x_j K_{kl}n^l\,dS\ , \qquad 2\cdot 10$$

we find that the vectors $\vec{P}$ and $\vec{S}$ are in fact the linear and angular momenta of our spacetime. And once again, just as with the Brill-Lindquist data, we can exploit the linearity of the momentum constraints to add freely multiple black hole solutions.

But we don't yet have a complete set of initial data for boosted spinning black holes. The Hamiltonian constraint still needs to be solved for ψ, which as a result of the non-vanishing extrinsic curvature can no longer be solved analytically. We do however still know the boundary conditions that are to be imposed on the solution—asymptotic flatness of the conformal factor at the outer boundary—which we assume once again satisfies behaviour of the form

$$\psi = 1 + \frac{k}{r}\ . \qquad 2\cdot 11$$

But the inner boundary requires a more subtle treatment. In the Brill-Lindquist data representing the Schwarzschild black hole we had a solution that is singular at $r = r_i$. By retaining the form of 2·11 here we retain the interpretation of the singularities as infinities in separate asymptotically flat regions and so generalise Brill-Lindquist initial data in what is known as *puncture* initial data. For multiple punctures we simply add terms of the form 2·8 for each black hole. Now in the Schwarzschild case we solved the momentum constraint by taking advantage of the time-independence, setting $\partial_t K = 0$ and vanishing lapse at the horizon. That option is no longer open to us, so without an inner boundary condition we will have to integrate right through to the other universe

and handle the singularities in ψ. This is done by separating out the singular piece, assumed to be of Brill-Lindquist form, from the rest of the solution u:

$$\psi = \psi_{BL} + u \, . \qquad \text{2·12}$$

Previously we saw how, for time-independent ψ_{BL}, the $\widetilde{D}^2 \psi$ term in the Hamiltonian constraint vanishes for a punctured manifold. We apply this to reduce the Hamiltonian constraint down to

$$\widetilde{D}^2 u = -\frac{1}{8\psi_{BL}^7}\widetilde{A}_{ij}\widetilde{A}^{ij}\left(1 + \frac{u}{\psi_{BL}}\right)^{-7} , \qquad \text{2·13}$$

and are left solving this in place of the Hamiltonian constraint in terms of ψ. Again, we start by considering boundary conditions and, again, asymptotic flatness means

$$u = 1 + \frac{k}{r} \, , \qquad \text{2·14}$$

but this time we have no need to worry about the inner boundary condition at the puncture as the divergent part now vanishes as you approach it. To see why, consider the Hamiltonian constraint 2·13. As we approach the puncture, ψ_{BL} diverges as $|\vec{r} - \vec{r}_i|^{-1}$, but the $\widetilde{A}_{ij}\widetilde{A}^{ij}$ component diverges as $|\vec{r} - \vec{r}_i|^{-6}$ (due to the $\frac{1}{r^3}$ in 2·8), and so

$$\lim_{r \to \infty} -\frac{1}{8\psi_{BL}^7}\widetilde{A}_{ij}\widetilde{A}^{ij} = \frac{|\vec{r} - \vec{r}_i|^7}{|\vec{r} - \vec{r}_i|^6} \, ,$$
$$= |\vec{r} - \vec{r}_i| \, . \qquad \text{2·15}$$

The Hamiltonian constraint then becomes $D^2 u = 0$. The existence, uniqueness, and regularity of solutions to this are discussed by Brandt and Brügmann in [25]. At this point we have in hand our complete set of initial data for a boosted spinning black hole. All that's left is to move this forwards in time using the evolution equations.

2.3. *Evolution*

Although it may be possible to move the punctures, the general sense in the community at the time was that this could not work, and early work relied on fixed punctures. The trouble, it was thought, with our solution in the previous section in 2·12 is that we

treat the conformal factor as two pieces, the regular time-dependent u and the singular time-independent ψ_{BL}. By relying on this time-independence in ψ_{BL} we are anchoring the puncture in place: by definition this piece is not changing with respect to time so we are limited to evolving static punctures. This is clearly insufficient if we want to simulate two inspiralling black holes. It should be noted that contrary to this some successful evolution schemes were proposed: notably [26] which applies the "Gamma-freezing" shift condition first introduced for excision evolutions in [27, 28] to punctures to permit motion of the horizon while keeping the inner asymptotic end of the puncture anchored, and [29] which makes use of dynamically adjusted co-moving coordinates to evolve two punctures for one orbital cycle.

Two methods were introduced to mitigate this issue, both of which involve abandoning the problematic splitting of the conformal factor. First, if we ensure that the domain is discretised such that at the start of the evolution the puncture lies between two grid points (see §2.5) then there is really no singularity in the computational domain and so the issue is avoided. The only concern would be in the calculation of derivatives across the region containing the would-be singularity due to how rapidly they change close to the puncture, and the accurate calculation of very large derivatives close to the puncture. We therefore define a new conformal factor here,

$$\phi = \ln \psi \,. \qquad\qquad\qquad 2 \cdot 16$$

The issue still remains of course—we will still be taking derivatives across a singularity—but a logarithmic divergence grows far slower than a reciprocal one. Nevertheless this does still produce large numerical errors, but these errors have been seen to be entirely contained within the interior of the black hole and so do not go on to destabilise the evolution [10].

The second option is to remove the singularity entirely through a simple redefinition of the conformal factor that forces it to vanish at the puncture,

$$\chi = \psi^{-4} \,. \qquad\qquad\qquad 2 \cdot 17$$

With the singular behaviour now absorbed into the vanishing χ at the puncture we have a completely regular conformal factor that can be evolved in its entirety. The puncture is free to move.

These initial data and evolution schemes have been, and continue to be, successfully used to simulate black hole spacetimes [30–32], though typically with a different slicing than that we've been using up to this point. We started this chapter with the introduction of the maximal slicing gauge condition, but in practice this involves solving an elliptic equation at each time step which is remarkably expensive. Instead, it is common to use the 1+log slicing, given by [26, 33],

$$\left(\partial_t - \beta^i \partial_i \right) \alpha = -2\alpha K \,, \qquad\qquad 2 \cdot 18$$

which shares the singularity-avoidance behaviour of maximal slicing without the computational expense. One intriguing feature of this slicing however is the behaviour of the isotropic coordinates as you approach the puncture. Rather than covering the interior of the black hole they double-cover the exterior, and so as we've seen our Bowen-York initial data for N black holes (FIG 2.1) represents $N + 1$ asymptotically flat regions connected by N wormholes—the singularity at $r = 0$ is then just a coordinate singularity rather than a curvature singularity. Upon evolution of the single black hole data it has been found [34] that while the geometry of the spacetime remains static as one would hope, gauge-dependent quantities are not; hence the coordinate system is dynamic. This dynamism persists only briefly before settling down to a new time-independent form, with a singularity that goes as $\frac{1}{\sqrt{r}}$ rather than $\frac{1}{r}$. It is worth emphasising that this structure *is* the stationary solution of the Schwarzschild spacetime in the 1+log gauge. Intuitively this can be pictured as the additional universes at the end of each wormhole detaching and the throat then asymptoting to a finite-area surface, or finite-radius cylinder in the embedding diagram representation of FIG 2.2. Note that this behaviour, like the wormhole behaviour before it, is not a feature of the interior topology but purely one of the coordinate system—the singularity at $r = 0$ remains just a coordinate singularity as we still have a punctured domain terminating at $r = 0$ to excise the curvature singularity. The observation of this coordinate behaviour has motivated the construction of initial data that skips over this transitionary period and begins directly with the relaxed gauge [34], labelled *trumpet* initial data for its form in FIG 2.2. There are no longer extraneous unphysical regions within each event horizon, and more importantly if the appropriate initial shift can be found there is hope that the evolution data would correspond to physical quan-
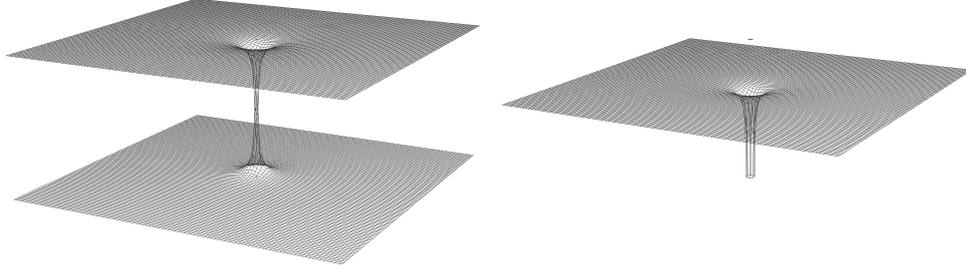
FIGURE 2.2   LEFT: *Interior coordinate topology of a 'wormhole' type puncture, represented as an embedding diagram of a spatial slice (inclination $\frac{\pi}{2}$) of the Kruskal-Szekeres maximal extension of the Schwarzschild solution. Two asymptotically flat external regions are connected by a throat of radius 2M, the event horizon [34].* RIGHT: *Interior coordinate topology of a 'trumpet' type puncture, represented as an embedding diagram of a spatial slice (inclination $\frac{\pi}{2}$) of the maximal extension of the Schwarzschild solution. The external region is asymptotically flat while internally it approaches a cylinder of radius $\frac{3M}{2}$ [34].*

tities, with minimal gauge evolution. It is hoped that this may in turn work to minimise the gauge motion of the puncture that occurs in the initial stages of each simulation [34].

All of this though, whether wormholes or trumpets, is built upon the assumption of conformal flatness. Bowen-York data is conformally flat, and reduces to the Schwarzschild spacetime—itself conformally flat in isotropic coordinates—in the static non-spinning case. It would be hoped, then, that in the static *spinning* case the Bowen-York data would reduce to the Kerr spacetime in some coordinate system, but it has been shown [35] that there exist no conformally flat slicings of the Kerr spacetime so this cannot be the case. This presents us with a problem: the Kerr solution is the only axisymmetric stationary black hole solution [36], so if the Bowen-York black hole we've derived, which is axisymmetric, *doesn't* reduce to Kerr the only resolution is that it must not be stationary. Indeed it turns out that what we have produced is a solution corresponding to Kerr plus some unphysical gravitational radiation, that we label *junk radiation*[2]. Thankfully it is sufficient to simply let this pass through the system; it doesn't destabilise the evolution and after its passage has no further impact on the data, though this does mean that the first $\sim 200M$ of data must be disregarded. While there is hope that the use of trumpet initial data may help to limit the resulting initial uncertainties in gravitational radiation [34], the conflict between conformally flat data and the Kerr spacetime remains the dominant source. Similarly, even though the boundary conditions in the BSSN formulation are

---

[2]The same is true of the boosted non-spinning case—the Bowen-York solution reduces to a boosted Schwarzschild black hole plus some spurious gravitational wave content

well-posed there still exist constraint violations that propagate inwards from the outer boundary of the domain as a wave [13]. In practice this is often 'solved' by pushing the boundary of the domain so far out as to be causally disconnected from the system being evolved in the centre.

In spite of these fundamental issues, our black hole spacetime can now be evolved through any one of a number of standard numerical methods for solving partial differential equations. But we still need an intuitive way to know what's going on in the simulation—we need to extract the dynamics of the black hole from the evolution of the spacetime as whole.

2.4. *Data Extraction*

Attributing mass and momentum to a black hole is nontrivial. Energy-momentum is an inherently non-local property of the spacetime [37]: for example, one option would be to calculate the ADM quantities we've already introduced in 2·9, but these are calculated at spatial infinity and refer the mass and angular momentum of the entire spacetime rather than the black hole itself. This wouldn't be a problem for an isolated, stationary Schwarzschild black hole but in the binary configurations we're interested in the spacetime will contain two black holes plus gravitational radiation. How much of the total mass and angular momentum 'belongs' to each of these components?

Typically [38–40] a black hole is defined by its event horizon, so we may be better served determining the dynamics of that. We can for example derive a measure of a black hole's mass the *irreducible mass* purely from the area A of its horizon,

$$M_{irr} = \sqrt{\frac{A}{16\pi}} \, , \qquad\qquad 2\cdot19$$

with the equivalent for cases with spin approximated by the Christodoulou mass,

$$M = \sqrt{M_{irr}^2 + \frac{S^2}{4M_{irr}^2}} \, . \qquad\qquad 2\cdot20$$

Unfortunately this isn't particularly practical for numerical simulations as the event horizon is defined as the chronological future of all points bounding those regions in the

spacetime not contained in the causal past of future null infinity[3], and so can be computed only *after* completing the simulation. We can however work with the *apparent* horizon, defined locally as the outermost closed two-dimensional surface on each slice whose outgoing null geodesics $\vec{L}$ orthogonal to the surface have zero expansion H everywhere. This surface lies within (or in the stationary case, coincides with) the event horizon [41]. If we let the 2D metric induced upon the surface by that of the slice be $h_{\mu\nu}$, the apparent horizon is then defined as the surface satisfying

$$H = -\frac{1}{2} h^{\mu\nu} \mathcal{L}_{\vec{L}} h_{\mu\nu} = 0 \ . \tag{2.21}$$

With a locally-defined horizon now available to us we can compute the corresponding mass from its area as above, and the angular momentum can be extracted following the procedure outlined in [42]. Recovering the dynamics of the black hole represents only half of the task we face though—ultimately we are interested here in the application to gravitational wave astronomy, so we now turn to the extraction of the wave content of our simulated spacetime.

    The gravitational wave signal is extracted at a finite radius from the binary in the form of the Weyl scalar Ψ_4, which we first need to reformulate in terms of our BSSN variables introduced in the last chapter. For completeness, and to make clear the conventions used, we recall the definition of the Weyl scalar,

$$\Psi_4 = -R_{\alpha\beta\gamma\delta} k^\alpha \bar{m}^\beta k^\gamma \bar{m}^\delta \ , \tag{2.22}$$

which is simply the projection of the full Riemann tensor $R_{\alpha\beta\gamma\delta}$ on to a coordinate basis formed of the null tetrad $l, k, m, \bar{m}$, where

$$-l \cdot k = m \cdot \bar{m} = 1 \ . \tag{2.23}$$

Following the Gram-Schmidt orthonormalisation procedure outlined in [43], we form the spatial triad

---

[3]Formally, the future event horizon in the manifold M is the boundary of the region $\mathcal{B} = M \setminus [M \cap \mathcal{J}^-(\mathcal{I}^+)]$, where $\mathcal{J}^-(x) = \{y \in M \,|\, y \prec x\}$ denotes the causal past of points x in M and $\mathcal{I}^+(x) = \{y \in M \,|\, y \ll x\}$ denotes the chronological future of points x in M.

$$u^i = [-y, x, 0] \, , \qquad\qquad\text{2·24}$$

$$v^i = [x, y, z] \, , \qquad\qquad\text{2·25}$$

$$w^i = g^{ia}\varepsilon_{abc}u^a v^b \, , \qquad\qquad\text{2·26}$$

with the orthonormal tetrad vectors given by

$$k^0 = \frac{1}{\sqrt{2}\alpha} \qquad k^i = \frac{1}{\sqrt{2}}\left(n^i - v^i\right) \, , \qquad\qquad\text{2·27}$$

$$l^0 = \frac{1}{\sqrt{2}\alpha} \qquad l^i = \frac{1}{\sqrt{2}}\left(n^i + v^i\right) \, , \qquad\qquad\text{2·28}$$

$$m^0 = 0 \qquad m^i = \frac{1}{\sqrt{2}}\left(u^i + i\,w^i\right) \, . \qquad\qquad\text{2·29}$$

Plugging these into 2·22 gives

$$\Psi_4 = -\frac{1}{4}\left(R_{abcd}v^a v^c - 2n^a R_{abcd}v^c + n^a n^c R_{abcd}\right)$$
$$\left(u^b - i\,w^b\right)\left(u^d - i\,w^d\right) \, , \qquad\qquad\text{2·30}$$

which when projected on to the slice becomes

$$\Psi_4 = -\frac{1}{4}\left(\mathcal{P}_a^p \mathcal{P}_b^q \mathcal{P}_c^r \mathcal{P}_d^s R_{pqrs}v^a v^c - 2\mathcal{P}_a^p \mathcal{P}_b^q \mathcal{P}_d^s n^s R_{pqrs}v^a + \mathcal{P}_b^q \mathcal{P}_d^s n^p n^r R_{pqrs}\right)$$
$$\left(u^b - i\,w^b\right)\left(u^d - i\,w^d\right) \, . \qquad\qquad\text{2·31}$$

Finally through the Gauss-Codazzi (1·3) and Codazzi-Mainardi (1·7) equations, together with 1·22, we can express the Weyl scalar purely in terms of our ADM variables:

$$\Psi_4 = -\frac{1}{4}\left[\left(\overset{3}{R}_{abcd} + K_{ac}K_{bd} - K_{ad}K_{bc}\right)v^a v^c - 2\left(D_b K_{ad} - D_a K_{bd}\right)v^a\right.$$
$$\left. + \left(\partial_0 K_{bc} - \beta^a \partial_a K_{bc} + K_{ab}\partial_c\beta^a + K_{bc}K_d^c + \frac{1}{\alpha}D_b D_d \alpha\right)\right]\left(u^b - i\,w^b\right)\left(u^d - i\,w^d\right) \, .$$
$$\text{2·32}$$

As $\mathit{\Psi}_4$ is of spin weight -2 we can project it on to the spin-weighted spherical harmonics
as

$$\mathit{\Psi}_4(t, \theta, \phi) = \sum_{\ell=2}^{\infty} \sum_{m=\ell}^{\ell} \mathit{\Psi}_{4,\ell m}(t) \; {}_{-2}Y_{\ell m}(\theta, \phi) \,, \qquad\qquad 2\cdot33$$

and evaluate it at some finite radius by integrating over a sphere, which also smooths
out noise in the numerical data.

This completes the journey from general relativity, through the numerical evolution
of black hole spacetimes, to the extraction of the associated gravitational radiation. We
have set up a theoretical framework that, when implemented, possesses all the features
necessary for the production of accurate simulated gravitational waves from black hole
binaries that are entirely consistent with analytical general relativity. It seems fitting
then to conclude this chapter with an outline of one such implementation: BAM.

2.5. *BAM*

First developed by Brügmann in 1997 [44], the Bifunctional Adaptive Mesh code, or
BAM , has seen active use over the past twenty years as numerical relativity has progressed
from fractions of orbits of Schwarzschild to full inspiral-merger-ringdown simulations
of precessing Kerr punctures, and as the study of gravitational waves has transformed
from a abstract theoretical endeavour to an observational science. Simulations gener-
ated by BAM were instrumental in this transformation, used in the search and analysis of
GW150914 along with those produced by other codes such as SpEC, LazEv, Maya, LEAN,
Llama, and others. These codes differ in a variety of ways further explored in CHAPTER 5,
but aside from SpEC all apply the puncture framework laid out in the preceding sections
to evolve a variety of spacetimes, sharing many features with BAM.

Leaving aside any physical interpretation, all we have at this stage is a set of partial
differential equations that we wish to solve numerically. The approach taken by BAM is
known as the *method of lines*, whereby the continuous spatial derivatives in our evolution
equations are approximated algebraically through fourth-order finite differences to leave a
dependence on time only. In this manner a partial differential equation is turned into a set
of ordinary differential equations, solved through fourth-order Runge-Kutta integration.

Within BAM the computational domain is discretised as a grid of N points at which
the equations are solved, separated by d. With any such discretisation scheme there is in-

evitably sometruncation error coming from the approximation of continuous derivatives with finite differences, an error that vanishes in the limit that $d \rightarrow 0$ (or equivalently resolution $\rightarrow \infty$). This will be of tolerable magnitude given some sufficiently fine spacing, a spacing that is determined by the length scale of the simulated system; for example in our black hole spacetimes it will be characterised by the short-wavelength features of our BSSN variables close to the puncture. Consequently a great deal of computational resources are wasted on the asymptotically flat space far from the puncture, which could be resolved perfectly well with much more widely spaced grid points. Ideally then we would use a changing grid spacing adapted to the changing length scale of the system: high resolution where there's fine features, low resolution where there's coarse features. This technique is known as *adaptive mesh refinement* and has been used to model everything from water flow around fishing nets [45] to the formation of galaxies [46].

The implementation of this within BAM consists of a series of nested Cartesian grids (or *boxes*) made up of N_l grid points of successively finer spacing, nested in the sense that the grid at each refinement level l is entirely covered by that of level $l - 1$. BAM uses the Berger-Oliger refinement scheme, in which the grid spacing d on each refinement level follows the scaling

$$d_l = \frac{d_0}{2^l} \qquad \{l \in \mathbb{Z}, l \geq 0\} \,, \qquad\qquad 2 \cdot 34$$

as does the duration between successive time steps. On the finest level the grids are centred on the black holes, one patch covering each, moving with the black hole. The extent of these grids is given by

$$D_l = \pm \frac{d_0}{2^{l+1}} \left(1 + N_l\right) \,, \qquad\qquad 2 \cdot 35$$

and as we move down the levels the grids grow in extent until a level is reached where the two would overlap (FIG 2.3). At this point they are instead replaced by a single fixed grid of minimal extent necessary to contain the two original grids, with the grid on the coarsest level covering the entire domain. Data is communicated between these levels of differing resolution through sixth-order polynomial interpolation. Strictly speaking this is not full adaptive mesh refinement, in the sense of using some accuracy criteria to dynamically adjust the size and shape of the refinement regions. In this treatment the
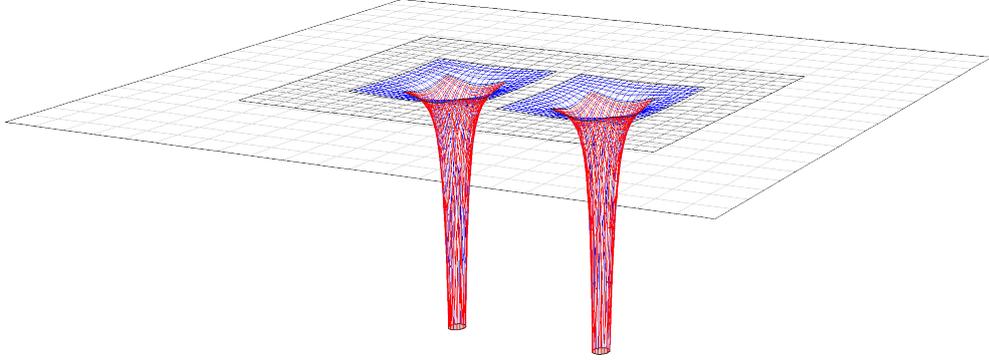
FIGURE 2.3   *Embedding diagram schematic of the grid structure used in BAM, shown here for an equal-mass binary. The adaptive mesh refinement procedure ensures that finely spaced grids are used close to the puncture while far more coarsely spaced grids are used far away. Here the grid spacings and extent scale by a factor of two between each level.*

boxes are moved to track the motion of the punctures and, since the black hole size and geometry changes little over the course of the evolution, this is found to be sufficient.

General numerics aside, in terms of specific numerical relativity choices BAM employs the BSSN formalism. For the production of simulations used for gravitational wave astronomy as in CHAPTER 4 black holes are evolved using the χ-variant of the moving puncture scheme described in §2.3, starting off with Bowen-York wormhole initial data and an initially vanishing "pre-collapsed" lapse of the form $\alpha = \psi_0^{-2}$ and shift vector $\beta^i = 0$. This is modified for later exploratory work presented in PART III, in which the ϕ-variant is used to evolve both wormhole and trumpet initial data as the need arises. In both cases the 1+log slicing is used, with the shift fixed by the $\widetilde{\Gamma}$-driver condition,

$$\partial_t^2 \beta^i = \frac{3}{4} \partial_t \widetilde{\Gamma}^i - \eta \, \partial_t \beta^i \ . \qquad\qquad 2 \cdot 36$$

where the η term is simply an *ad hoc* addition that damps strong oscillations in the shift. In all simulations produced for this thesis we set $\eta = 0$ following the observation in [34] that doing so minimises the coordinate drift between the use of wormhole- or trumpet-form 1+log data. 2·36 describes a common technique to solve an otherwise computation-

ally expensive shift condition (cf. the Gamma-freezing condition $\partial_t \widetilde{\Gamma}^i = 0$ introduced in [28]) by setting up β^i as an initial value problem.

It is the shift vector that provides our first glimpse into the dynamic behaviour of the puncture in these simulations as 1·35 implies that where $\psi \to \infty$ (that is, at the puncture) the velocity $\partial_t \vec{x} = -\vec{\beta}$ [47], though recall that BAM's staggered grid configuration means that we will have to interpolate on to the puncture to extract this speed. Using this, and by first estimating that at the current time step the puncture is located in the same place as at the previous, BAM is able to iteratively improve its estimate of the puncture's location $x_{current}$ by considering

$$
\begin{aligned}
x_{current} &= x_{previous} + v\,dt\,, \\
&= x_{previous} + \frac{-\beta_{current}(x_{current}) - \beta_{previous}(x_{previous})}{2}\,dt\,.
\end{aligned}
$$

<div align="right">2·37</div>

This goes through three iterations to arrive at a final estimate of the current location. The rest of the puncture dynamics data is extracted from the apparent horizon following the procedure outlined in §2.4, which is itself located using the fast flow algorithm (introduced in [48]) to find the surface that satisfies 2·21. This involves placing a spherical 'trial surface' of points x^i well outside the expected location of the horizon, then iteratively drawing each point inwards by a distance proportional to $H(x^i)$. In this way the trial surface will continue to contract until $H = 0$ is satisfied at all points (to within a given tolerance), like a vacuum-packed fish.

The gravitational waveform is extracted as the projection of Ψ_4 interpolated on to a sphere of finite radius (the *extraction radius*), as outlined in §2.4. The interpolation is performed using fifth-order polynomials and the integration over this sphere is approximated using fourth-order Runge-Kutta. But whereas our adaptive mesh refinement scheme is perfectly suited to the dynamics of the black hole, it poses something of a problem for the radiation. The wavelength of the gravitational wave signal stays constant with distance and so requires constant radial resolution, but we also need increasing accuracy to distinguish it from the background as the amplitude falls off as $\frac{1}{r}$. The solution in BAM is to use two different values for the number of points for the boxes at each level: one for tracking the punctures and one for wave extraction.

Waveform Modelling

3

The previous chapter has taken us from the setting up of the initial state of a Schwarzschild black hole all the way through to the extraction of a simulated gravitational wave signal in BAM. The aim of all this is, as far as the scope of this thesis is concerned, to use these simulated waveforms in the search for and analysis of real signals from real black hole binaries picked up by gravitational wave detectors. With the groundwork now laid out and the ability to simulate generic black hole spacetimes, the next step is to populate the space of possible parameters of a black hole binary with representative numerical waveforms. Unfortunately with eight independent parameters (m_1, m_2, S_1, S_2) a satisfactory sampling would require on the order of millions of simulations. Even utilising parallel processing on high-performance computing clusters each of these takes weeks to run for sufficient duration, so this is clearly not feasible. Instead, a number of models have been developed that, for a given set of parameters of a binary, return a rapid approximation to the true waveform, an approximation that is sufficiently accurate to successfully identify signals in the detector data and estimate source parameters for the signal-to-noise ratios expected in current gravitational wave detectors. There are a number of different approaches to this that we outline in this chapter, all of which typically model the decomposition of the signal into spherical harmonics as outlined in 2·33. In practice however it is not the Weyl scalar but the strain, defined [10]

$$ h(t) = \lim_{r \to \infty} \int_0^t \int_0^{t'} \Psi_4(t'') \, dt'' dt' \, , \qquad\qquad 3\cdot1 $$

that's modelled as it is this quantity that the interferometers measure. The coefficients $h_{\ell m}$ are formally the multipole moments of h, but are far more commonly referred to in the literature as *modes*, and it is the amplitude A and phase ϕ of these modes that are modelled, where

$$h_{\ell m}(t) = A_{\ell m}(t)e^{-i\phi_{\ell m}(t)} \,.$$

<div align="right">3·2</div>

The signal is then expressed as

$$h(t,\theta,\phi) = \sum_{\ell=2}^{\infty}\sum_{m=-\ell}^{\ell} h_{\ell m}(t) \,_{-2}Y_{\ell m}(\theta,\phi) \,.$$

<div align="right">3·3</div>

In an equal-mass binary whose component black holes are non-spinning, or have their spins (anti)aligned with the orbital plane, the ($\ell = 2, |m| = 2$) modes dominate the spectrum and as such it is only more recently that the higher $\ell = |m|$ modes and the subdominant $\ell \neq |m|$ modes have been modelled.

3.1.   *The Post-Newtonian Formalism*

The earliest efforts towards modelling gravitational wave emission from a compact binary were simple perturbative expansions in the low-velocity weak-stress limit: the *Post-Newtonian* (PN) expansion. This technique dates back to the birth of general relativity itself having first seen use in the work of de Sitter, Lorentz, and Droste in the mid-1910s computing the equations of motion for an N-body system [49,50]. The derivation of the radiation-reaction terms in the expansion came later, and with orbital phase proportional to gravitational wave phase [51] this produced the first model for the phase of the waveform. These results have been pushed to higher and higher orders in the expansion over the past century, and different approaches to the expansion have been explored over the years (see [52] for a thorough summary and comparison). Of course, the closer we move towards merger in the evolution of the binary the further we depart from the low-velocity weak-stress approximation, and the more inaccurate our PN results become.

3.2. *Black Hole Perturbation Theory*

At the other end of the scale we know that the final state of an inspiralling will be a single isolated remnant black hole whose gravitational wave output can be modelled trivially: there won't be any. But this state doesn't occur instantaneously upon merger. The merger of the two black holes can be defined as the moment that a common apparent horizon forms around them. This common horizon will be highly distorted, far from the neat spherical symmetry or axisymmetry of an isolated black hole horizon, but as it is now simply a single perturbed black hole it can be tackled with standard perturbation theory techniques. Comprehensive reviews of black hole perturbation theory are available in [53, 54], but for the purposes of this brief outline it will suffice to say that for a small perturbation a Kerr black hole will relax back to its stationary state, radiating away the perturbation in the form of gravitational waves. This process is known as *ringdown*, highlighting the analogous behaviour of a struck bell. The equivalent of a bell's characteristic overtones, the *quasinormal modes*, can be calculated within black hole perturbation theory and supply the final stages of a waveform model.

3.3. *The Effective One-Body Formalism*

The first complete models, of inspiral through merger and into ringdown, used the *effective one-body* (EOB) formalism. To model the inspiral, the PN expressions for the dynamics of two masses are resummed to represent the dynamics of a test particle on a deformed Kerr background, as are the radiation-reaction expressions and the gravitational wave strain. This is then stitched together with results from black hole perturbation theory for the post-merger portion of the waveform using a simple step function. As may be expected these models perform poorly around merger, where neither formalism is valid: PN theory is based upon the assumption of a weak gravitational field, black hole perturbation theory assumes small perturbations, but merger is both in the strong field and involves a highly deformed black hole. For this reason later models in the EOB family have included a number of free parameters that can be tuned to adjust the merger portion of the waveform based on numerical waveforms. Recent achievements in the EOB family include the incorporation of precession effects and higher modes. For a good summary see [55].

Phenomenological Modelling

Both the PN and EOB class of models are based on an understanding of the underlying
physics of the inspiralling system: they approximate the dynamics and from that pull
out the behaviour of the waveform. The *phenomenological* (or more commonly just *phe-
nom*) models, as their name suggests, cut out the middle man and directly approximate
the waveform itself, initially with a functional form inspired by PN and perturbation
theory results though more recently hybridised from numerical simulations and PN ap-
proximants. Although these ingredients themselves rely on an accurate understanding
of the physics involved, the key point is that this is completely divorced from the result-
ing model. A defining feature of phenomenological models is the use of closed-form
analytic expressions; they are therefore significantly cheaper computationally than EOB
models, but are historically much more dependent on numerical waveforms—and there-
fore much more dependent on them being accurate. In recent years however phenom
and EOB models alike have been tuned using numerical waveforms and so are both in
some sense phenomenological.

The first phenomenological model used a piecewise construction of PN approximants
for the inspiral and a set of non-spinning numerical simulations of mass ratio $1 \leq q \leq 4$
binaries for merger-ringdown, with some matching function in between constructed so
as to minimise the integrated phase and frequency difference between the two. Nested
fits were then made to the amplitude and phase data of these hybrid waveforms, first
fitting a piecewise ansatz to the waveforms, then making polynomial fits to those fitting
coefficients to form the final model. This is an approach that has remained unchanged
even as the models have grown to include increasingly complex physics.

If for example the total spin angular momentum of the binary is not (anti)aligned
with its orbital angular momentum, the orbital plane will steadily precess around the total
angular momentum. As gravitational radiation is beamed predominantly along the direc-
tion of the orbital angular momentum [51] this precession of the binary leads to modula-
tions in the gravitational wave amplitude received by a static observer. Specifically, the
$(2, \pm2)$ modes no longer dominate the spectrum with power shunted into the previously
subdominant harmonics. In order to capture this behaviour precessing phenomenologi-
cal models define some frame in which these precession effects are minimised by noting
that an observer located directly above the orbital plane will experience the maximal ra-
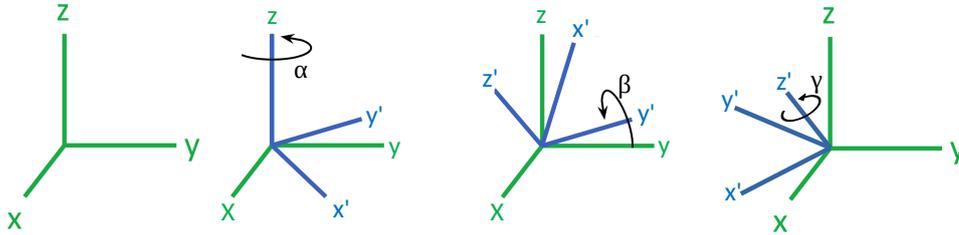
FIGURE 3.1   *Any rotation in 3D space can be represented by three angles, labelled here α, β, γ. The set of green axes represent the initial orientation, and the blue represent their latest orientation at each step.*

diated power from the binary. If this observer were to remain in a frame fixed in this position with respect to the binary, tracking the precession dynamics, they will observe significantly less modulation of the $(2, \pm 2)$ mode and the signal will resemble that from a non-precessing binary. In the literature this frame is called variously the coprecessing frame, the quadrupole-aligned frame, or is said to be aligned to the optimal emission direction [56–58]. The coprecessing frame can be located by finding at every moment that frame in which power in the $(2, \pm 2)$ mode is maximised. It is useful therefore to define the rotation that will take us into this coprecessing frame. We define this using the Euler angle notation (for its equivalent using the quaternion notation, see [59]), where by Euler's rotation theorem any rotation in three-dimensional space can be described using just three angles. In fact only two Euler angles are necessary to rotate a binary into the coprecessing frame, with the third acting as simply a final phase shift, a rotation of the black holes around their orbits. The precise definition of these rotations is a matter of convention, both in the labelling of the angles and in the axes used, but in this work we define them as follows (illustrated in FIG 3.1):

1. A rotation of α about the initial z-axis,
2. A rotation of β about the newly-formed y-axis,
3. A rotation of γ about the newly-formed z-axis.

These *precession angles* allow us to effectively remove the precession effects from a precessing signal, but more importantly for waveform modelling can be used to 'twist up' a non-precessing waveform to generate its precessing counterparts. This was precisely the insight that lead to the development of the first precessing waveform model in the frequency domain, PhenomP in 2014 [60], which has since been refined through the tuning of its precession angle ansätze to a set of 80 numerical waveforms [61].

Unequal component masses also serve to amplify the subdominant modes as odd ℓ modes are no longer suppressed when $q > 1$ [62], so in progressing to higher mass ratios it is increasingly important that our waveform models include more than just the $(2, \pm 2)$ mode. In 2018 London et al. presented the first higher-mode model of gravitational waves, `PhenomHM` [63]. This was achieved in a qualitatively similar manner to `PhenomP` only rather than turning a non-precessing waveform into a precessing one, it turns the $(2, \pm 2)$ mode into the higher modes through simple piecewise linear mappings,

$$
\begin{aligned}
A_{\ell m}(f) &\rightarrow |\beta_{\ell m}(f)| A_{22}(f_{22}) \,, \\
\phi_{\ell m}(f) &\rightarrow \frac{\phi_{22}(f_{22})}{f'_{22}(f)} + \Delta_{\ell m} \,.
\end{aligned}
\tag{3.4}
$$

involving the phase and frequency of the $(2, 2)$ mode, ϕ_{22} and f_{22}. The mappings $\beta_{\ell m}(f)$ and $\Delta_{\ell m}$ are determined by black hole perturbation theory during the ringdown portion of the waveform, and during inspiral by two facets of PN theory: first that the frequency of each mode (ℓ, m) is approximately m times the orbital frequency, and second that phase remains approximately stationary. These two regimes are joined by linear interpolation.

3.5. *Gravitational Self-Force*

Separate to the above approaches, which are used in the comparable-mass ($q \lesssim 10$) regime, are perturbative calculations for extreme mass ratio ($q \gtrsim 10^5$) inspirals in which the smaller component black hole is modelled as a point mass on the background of the larger. Perturbing this background black hole spacetime to first-order in the mass-ratio takes into account physical effects which arise from the self-interactions between a mass and its own gravitational field—the *gravitational self-force*, which is the dominant cause of inspiral in the binary. This is calculated by evaluating derivatives of the metric perturbation at the location of the particle, which is singular and so requires regularisation of the perturbation to remove the singularity.

Such first-order calculations are sufficient for the tracking of the smaller component over a few gravitational wave cycles, but the neglected higher-order terms accumulate as time goes on and introduce dephasing errors. Second-order perturbation theory however is far more complex than first-order for non-linear field theories, and extensive efforts

are currently underway to extend self-force calculations in hopes of achieving sufficient phase accuracy for LISA observations. A thorough introduction to self-force as it relates to waveform modelling can be found in [64].

A further issue that will need to be addressed in this regard is the issue of *resonances*, where (for certain configurations) the inspiral passes through points where the quotient of two characteristic orbital frequencies becomes rational and produces a 'jump' in phase that can't be captured by current self-force calculations [65]. The use of a waveform model without such discontinuities in the analysis of a signal that features them will lead to a non-negligible bias in the science output of these space-based observatories [66]. Accurate waveforms will be essential for them just as they have been for the current generation of ground-based observatories, needed both in searches for gravitational wave signals and for the inference of their source, and it is the application of these models to such things that we address in the next section.

3.6. *Searches*

However you have arrived at your model for the gravitational waveform, whichever formalism has been used in its construction, with one in hand we are in a position to generate a bank of template waveforms that cover the vast parameter space in a way that would have been entirely unfeasible with computationally expensive numerical simulations. This template bank can then be compared against the data stream from a given gravitational wave detector to see if it contains any signals of a morphology suggestive of it having been generated by an inspiralling binary. It is this procedure we now address, outlining how we perform and quantify such a comparison with data dominated by noise, and how we can apply our template bank to extract probable parameters of the source binary.

The output of each detector is a continuous stream of noise potentially peppered with gravitational wave signals hidden amongst that noise. These signals, if they are sufficiently loud, are picked out by a *matched filter* search which finds, for some data D, the function F that maximises the inner product C between them:

$$C(\tau) = \int_{-\infty}^{\infty} F(t + \tau)D(t)dt \, ,$$

3·5

for the cross-correlation variable τ. Taking the Fourier transformation we find that in the frequency domain this becomes

$$C(\tau) = 2\,\mathrm{Re}\int_0^\infty \widetilde{F}^*(f)\widetilde{D}(f)e^{-2\pi i f\tau}df\;.$$

<div align="right">3·6</div>

This is more convenient than working in the time domain because it allows the noise in the data to be characterised by its power spectral density S_n^d alone, if we assume it to be stationary and Gaussian. For a signal h buried in this noise the optimal filter—that which maximises the signal-to-noise ratio—would by definition be that same signal divided by the noise,

$$\widetilde{F}_{opt}(f) = \frac{\widetilde{h}(f)}{S_n^d(f)}\;,$$

<div align="right">3·7</div>

but of course it is impossible to know the shape of the signal before it arrives—the general form is known well enough but as we've seen it depends heavily on the mass ratio, spin configuration, and inclination. To find the optimal filter we therefore take our bank of model waveforms and run each of them through 3·6 to see which gives the highest correlation. Replacing h with the general model G, representing any waveform within the template bank, the optimal filter is given by

$$\widetilde{F}_{opt}(f) = \frac{\widetilde{G}(f)}{S_n^d(f)}\;,$$

<div align="right">3·8</div>

which substituted into 3·6 gives

$$C(\tau) = 2\,\mathrm{Re}\int_0^\infty \frac{\widetilde{G}^*(f)}{S_n^d(f)}\widetilde{D}(f)e^{-2\pi i f\tau}df\;.$$

<div align="right">3·9</div>

We can simplify this by defining the *noise-weighted inner product* between two waveforms A and B,

$$\langle A|B\rangle = 2\,\mathrm{Re}\int_0^\infty \frac{\widetilde{A}^*(f)\widetilde{B}(f)}{S_n^d(f)}df\;,$$

<div align="right">3·10</div>

so that the correlation can now be written as

$$C(\tau) = \langle Ge^{-2\pi i f\tau}|D\rangle\;.$$

<div align="right">3·11</div>

This can be simplified by using merger time as our cross-correlation variable τ, which we are free to set as $t_{merge} = 0$ to leave

$$C(\tau) = \langle G | D \rangle \ . \qquad\qquad 3 \cdot 12$$

This noise-weighted inner product provides a rigorous way to quantify the agreement between different waveforms. We first normalise the waveforms and shift their parameters p such that they are optimally aligned with respect to one another, and so define the *match* between two waveforms as

$$\mathcal{M}(h, G) \equiv \max_{p} \frac{\langle h(\phi, t) | G \rangle}{\langle h | h \rangle \langle G | G \rangle} \ , \qquad\qquad 3 \cdot 13$$

where 0 represents complete disagreement and 1 complete identity between h and G. It is also common in the literature to work in terms of the *mismatch*, given by simply $1 - \mathcal{M}$. The set of parameters p over which to optimise is dependent on context: in CHAPTER 5 when comparing two precessing numerical waveforms to one another we choose phase, time, and polarisation. For the $(2, 2)$ mode of a non-precessing binary however we can exploit the degeneracy between phase and polarisation (see [67]) to use a simple phase shift as a proxy for varying polarisation. In the precessing case the subdominant modes, each of which have their own associated phases, can no longer be neglected and so the degeneracy is broken, which is why the precessing matches need to be optimised over polarisation explicitly. In addition to this by definition a precessing binary is one with a time-varying inclination, which we have seen in §3.4 modulates the observed power in each mode, so this relative inclination between the two signals is one further parameter over which to optimise. In CHAPTER 5 this is instead used as an independent variable when computing matches. If, rather than comparing two known waveforms, we have some detector data containing an astrophysical signal then this optimisation is equivalent to measuring each of the parameters—we are finding the parameter values that best match the signal. This is the process of *parameter estimation*, which we briefly summarise below for completeness.

3.7. *Parameter Estimation*

Having identified a candidate signal in the detector data we would like to put it to work. Just like traditional electromagnetic astronomy, by analysing a signal we hope to extract information about the astrophysical object that generated it. We've seen that the waveform is highly dependent on the parameters of the binary from which it originated, so it should in turn be possible to infer those parameters from the detected waveform. This is the motivation behind parameter estimation (PE) efforts in the gravitational wave community. PE is fundamentally an exercise in Bayesian inference, which at its heart has that the probability of event A given that event B has occurred is proportional to the probability of B given A multiplied by the probability of A:

$$P(A|B) \propto P(B|A)P(A) \,. \qquad 3\cdot14$$

For example, you pick up a coin from the street. Either this is a fair coin (that is, it has a heads side and a tails side) or it is unfair (it has two heads sides). $3\cdot14$ tells us that the probability that it is fair, given you flipped it and got a heads, is proportional to the probability of getting a heads if it really had been fair multiplied by how likely you would have thought it yesterday that a coin you find on the street is fair. You now have more information, so the next time you flip the coin you are able to update how likely you think it is that a coin you find on the street is fair. Each time you flip the coin provides more information that can serve to update your prior assumption about how likely that is—if you flipped it 500 times and it came up heads every time, you might well start to think that finding an unfair coin on the street is more likely than you previously imagined.

PE is the generalisation of this to a far larger parameter space with far more possible states. Rather than the two options 'fair' or 'unfair' for the coin in our example above we are faced with determining the which of the millions of possible parameter combinations contained within our template bank is the likely state of a detected signal. It is in essence, for a signal s with parameters p and signal model h, the calculation of the quantity

$$P\left(p|s,h\right) \propto P\left(s|p,h\right) P\left(p|h\right) \,. \qquad 3\cdot15$$

These three probability distributions are known as the *posterior*, *likelihood*, and *prior*. The posterior is the probability that the signal s has the parameters p given you've detected

this signal s and you have the model h for signals, the likelihood is the probability of detecting the signal s if it really did have parameters p according to your model h, and the prior is how likely your model h says you are to get those parameters p before taking into account that you've received the signal s.

In practice the priors are informed by models of the likely black hole population of the universe derived from studies of star formation, models of active galactic nuclei, electromagnetic observations of black holes, and a host of other areas including most recently previous gravitational wave observations. In all but the most trivial scenarios the posterior cannot be written down analytically, so typically within the LVK the posterior is computed through nested sampling or the Markov Chain Monte Carlo method (see [68] for summaries of these and other sampling algorithms).

This completes the journey from abstract geometry to tangible, practical tools that can be applied to real data to probe astrophysical phenomena. The Einstein field equations have been turned into a tractable computational problem, with black holes as a key result of general relativity represented accordingly; we've explored one particular implementation of this framework, BAM; and have motivated the production of numerical simulations of black hole binaries. It is with these simulations that we begin the next part of this thesis.

Part Two

CONTEMPORARY NUMERICAL SIMULATIONS

The BAM Catalogue

4

We have explored in PART I a little of the machinery behind the successful solution of the Einstein equations for a black hole binary, and discussed the role that these simulations play within the field of gravitational wave astronomy. We have seen how the gravitational waveforms they produce have been used in the development of waveform models that were an integral part of the historic first detection of gravitational waves and continue to be applied to every detection to date.

Both the EOB and phenomenological models rely on numerical waveforms for accuracy around merger: for example, at the time of the first gravitational wave detections SEOBNRv2 was calibrated to 30 aligned-spin and 8 non-spinning simulations up to $q = 8$ [69], while PhenomD was tuned using 19 simulations that uniformly cover the $\frac{1}{q}$ aligned spin parameter space [70]. Both of these have been used extensively during the first and second observing runs of the LIGO detector and there are now increasingly accurate precessing, higher-mode models in both the SEOBNR and Phenom families [60, 63, 71–74], and also now surrogate models constructed entirely from NR waveforms [75–78]. Numerical waveforms are thus a vital step in the 'production line' of scientific output in gravitational wave astronomy.

A number of groups have in recent years released catalogues of simulated waveforms for just this purpose. The largest of these consists of, at the time of writing, 2028 waveforms produced by the SXS collaboration using the SpEC code [79], followed by 777 from RIT using LazEv [31], and 452 from Georgia Tech using Maya [32]. These catalogues are summarised in TABLE 4.1. Given the computational expense the high mass ratio regime is understandably under-explored compared to the equal-mass regime—naïvely

one may assume that this doesn't represent any particular cause for concern seeing as current detections almost all lie well within the approximately equal-mass regime [80]. However for the vast majority of these, despite having most support at equal mass, parameter estimation results produce 90% credible intervals extending as high as $q = 8$. And with 50 published detections to date including a $q = 9$ event [81], and a further 80 [7] expected in the fourth observing run the prospect of detecting high mass ratio event grows ever more likely, with a great need for simulations in this region.

Catalogue	Number	q	χ_1	χ_2
Cardiff	80	1–8	0.000–0.000	0.200–0.800
SXS	2028	1–16	0.000–0.998	0.000–0.496
RIT	777	1–15	0.000–0.951	0.000–0.953
GT	452	1–15	0.000–0.809	0.000–0.809

TABLE 4.1   *Parameter space coverage offered by four catalogues used in the production of waveform models, where $\chi_i \equiv S_i/M_i^2$ defines the dimensionless spin magnitude.*

More broadly, rather than any systematic sampling of the parameter space simulations have been produced organically as the field and the capabilities of the codes grows. This has left us with a highly non-uniform coverage of the parameter space and it is in an attempt to address this, as well as the dearth of simulations in the high mass ratio high spin regime, that the Gravity Exploration Institute at Cardiff University will soon publish its first public catalogue of simulations. With 80 single-spin black hole binary configurations it is a somewhat more modest affair than the other catalogues currently available, but importantly it represents a systematic effort to cover a subset of the parameter space that includes high mass ratios and high spins. The results presented in this chapter are a summary of those due to appear in an upcoming paper [82].

High mass ratio simulations pose a significant challenge for all current numerical relativity codes, so producing the large number of simulations at $q = 8$, of a length suitable for accurate waveform modelling, was a nontrivial task. Addressing this challenge at significantly higher mass ratios is the motivation behind PART III of this thesis. Here we iterate over four mass ratios $q = \{1, 2, 4, 8\}$, four spin magnitudes $\chi_2 = \{0.2, 0.4, 0.6, 0.8\}$, and five orientations of the spin with respect to the Newtonian orbital angular momentum $\theta = \{30°, 60°, 90°, 120°, 150°\}$.

This is the product of the continued effort of a number of people over the past four years, of whom the author is just one among many, with the motivation of providing the numerical waveforms necessary for the development and calibration of precessing waveform models. In this chapter we briefly revisit some of the material covered in §2.5 regarding the setup of BAM simulations, expanding upon those points of particular relevance to the simulations featured in this catalogue, before providing an overview of catalogue and finally discussing a number of accuracy considerations. The work presented herein is the result of a collaboration between Edward Fauchon-Jones, Eleanor Hamilton, Mark Hannam, Charlie Hoy, Chinmay Kalaghatgi, Lionel London, Jonathan Thompson, Shrobana Ghosh, Sebastian Khan, Panagiota Kolitsidou, Alex Vañó-Viñuales, and the author.

4.1. *Numerical Setup*

All the simulations performed to produce the catalogue presented here were performed using BAM (see §2.5 for an in-depth introduction to the code). Black holes are evolved using the χ-variant of the moving puncture scheme, represented with Bowen-York wormhole initial data together with an initially vanishing "pre-collapsed" lapse of the form $\alpha = \psi_0^{-2}$ and shift vector $\beta^i = 0$. Gauge choices take the form of the 1+log slicing condition and the $\widetilde{\Gamma}$-driver conditions.

Two restrictions are placed on the particular setup of the numerical grid structure for the sake of accuracy, based purely on prior experience producing BAM simulations. First, the width of the finest box around a black hole should be 1.2-1.5× the maximum diameter of the (pre-merger) apparent horizon, which is known from a lower resolution version of the simulation. Second, on the level that the gravitational wave signal is extracted there should be at least ten grid points per wavelength of the $(4, 4)$ mode. As the wavelength decreases as you approach merger this requires foreknowledge of the smallest wavelength that will be reached in the simulation, provided in the form of an estimate from the ringdown frequencies f_{RD}, approximated here by the PhenomD waveform model. The grid spacing on the wave extraction level should then be $d_l = \frac{1}{20f_{RD}}$. In cases where it's not possible to satisfy both of these requirements, box size and extraction-level grid spacing are balanced to minimise the extent that either is violated.

The desired spin orientations of the black holes are achieved through the iterative algorithm presented in [30]. Briefly, the orientation is provided initially at large sep-

aration and the system is evolved via the EOB equations of motion up to the desired reference orbital frequency Ω at which the simulation will begin (the *starting frequency*). Of course by doing this we end up with a slightly different spin configuration at the start of the simulation than we were aiming for because the spin will have evolved in the intervening time, so the difference between the actual and desired spin directions is noted and the initial orientation is rotated by that amount. This is again evolved forwards to the starting frequency using the EOB equations and again the discrepancy is checked and the initial guess rotated. This is repeated until initial parameters are obtained that produce the desired configuration at the starting frequency to a tolerance of 1% in Ω.

The initial parameters are then further manually iterated to lower the eccentricity below a threshold set at 0.002 over the first 1000M of simulation time, motivated in part by [83]. We estimate this following [84], fitting the puncture separation between simulation time 200–700 M using a quadratic function and taking the eccentricity to be the maximum difference between the fit and the data in this range.[1] Small perturbations are introduced to the magnitude of the black holes' linear momenta, typically for the simulations in this catalogue on the order of 0.1–0.8%, until the eccentricity is below the threshold. If this cannot be achieved through perturbation of the linear momenta alone, the radial momenta are decreased by 25–75%. For the sake of speed these eccentricity-reduction runs are performed at lower resolution than will be used in the final product, typically with grid spacings 1.5× larger and wall-clock speeds three times faster. As a result of this differing numerical accuracy the eccentricity of the final simulation may well differ from that of the eccentricity-reduction run, in some cases rising above the 0.002 limit we have set. This has been seen in a number of the simulations within this catalogue, but in each case the deviation is considered small enough to be acceptable. For a summary of these eccentricities see TABLE 4.2.

4.2. *Description of the Catalogue*

The 80 simulations contained within the catalogue can be separated into 20 simulations at each of four mass ratios $q = \{1, 2, 4, 8\}$. Each subset of 20 can be further divided into five simulations at each of the four spin magnitudes $\chi_2 = \{0.2, 0.4, 0.6, 0.8\}$. All of

---

[1]The full method described in [84] requires the merger time, and is used to estimate the eccentricity of the final simulation.

these simulations are single-spin, with vanishing spin on the smaller black hole as the contribution to the overall precession dynamics of spin-spin couplings is expected to be subdominant [85]. These five simulations vary the orientation of the spin; specifically, the angle between the spin vector on the larger black hole and the Newtonian orbital angular momentum takes on the values $\theta = \arccos\left(\hat{\vec{L}}_N \cdot \hat{\vec{\chi}}_2\right) = \{30°, 60°, 90°, 120°, 150°\}$ to within a tolerance of 1°, negligible for the signal strengths expected in current-generation detectors.

The simulations are required to be of sufficient duration that they can be connected to the PN waveform models of early inspiral while it retains its validity. Computational expense aside, the main obstacle faced by longer-duration simulations is that numerical error in the phase of the binary is cumulative, improved only through the use of finer resolutions. We are then faced with determining some balance between a resolution fine enough to keep this dephasing error within acceptable limits, yet coarse enough that the computational expense doesn't become prohibitive. Through previous experience with $q = 18$ simulations we settle upon a resolution that typically lies within the range $0.01m_i$–$0.02m_i$, where m_i refers to the mass of either component black hole $i = 1, 2$, and a duration of approximately $2000M$, equivalent to roughly 10 orbits. Initially it had been planned that each of the simulations would begin at the same starting frequency, $M\Omega = 0.023$, but spin-orbit coupling can, depending on the magnitude and orientation of the spins on the black holes, hasten or indeed delay the onset of the merger phase. For the simulations in our catalogue this is particularly pronounced in the higher mass ratio configurations with spins inclined 30° to the orbital plane, which were seen to have merger times over $3000M$ when starting at the prescribed starting frequency. The cumulative phase error at such a duration would introduce unacceptable inaccuracies in the resulting waveform, and the higher resolutions required to counteract it would further lengthen the 3 month production timescale for this single simulation, so instead we simply begin at a higher frequency.

The key features of these simulations are summarised in TABLE 4.2 and TABLE 4.3 with further details available in [82], where the effective and precessing spin parameters are defined respectively as

$$\chi_{\text{eff}} \equiv \frac{m_1}{M}\left(\vec{\chi}_1 \cdot \hat{\vec{L}}\right) + \frac{m_2}{M}\left(\vec{\chi}_2 \cdot \hat{\vec{L}}\right), \qquad\qquad 4\cdot1$$

$$\chi_p \equiv \frac{1}{2m_2^2} \max\left(\frac{3+4q}{4+3q} \left(\vec{\chi}_1 - \left(\vec{\chi}_1 \cdot \hat{\vec{L}} \right) \hat{\vec{L}} \right), \; \vec{\chi}_2 - \left(\vec{\chi}_2 \cdot \hat{\vec{L}} \right) \hat{\vec{L}} \right).$$ 4·2

4.3. *Accuracy Analysis*

Four of the 80 simulations have been analysed by Eleanor Hamilton to estimate the accuracy of the resulting waveforms. Such an estimate is essential given their intended use as it provides an upper bound on the accuracy of any waveform model developed using this data. The four simulations are those with the configurations

$$(q, \chi_2, \theta) = \begin{cases} (4, 0.4, 60°) \\ (4, 0.8, 120°) \\ (8, 0.4, 30°) \\ (8, 0.4, 150°) \end{cases}.$$ 4·3

Errors in the numerical waveforms used by the gravitational wave community are dominated by resolution and wave extraction—specifically, the truncation error due to the use of a finite grid spacing, and the error due to extraction of the gravitational wave at some finite radius. It is in these two areas that focus is directed in this section, employing a similar approach to quantify the error in each case. In particular we focus on the match, the amplitude and the phase of the $(2, 2)$ mode in the quadrupole-aligned frame, and the two precession angles α and β. The key tool employed to quantify these errors is *convergence testing*.

Take two runs of the same simulations with different grid spacings d_1 and d_2 that report the solutions $u_1(t)$ and $u_2(t)$. The dominant error term in any simulation goes as d^n, which is to say the 'true' solution is given by

$$u(t) = u_1(t) + k(t)d_1^n + \mathcal{O}\left(d_1^{n+1} \right),$$ 4·4

with some function $k(t)$. The error between our numerical solution and this true solution is then

$$u(t) - u_1(t) = k(t)d_1^n + \mathcal{O}\left(d_1^{n+1} \right),$$ 4·5

q	χ_2	$\theta\degree$	χ_{eff}	χ_p	e $(\times 10^{-3})$	$M\omega_{orb}$	N_{orb}
1	0.2	30.0	0.087	0.100	1.51	0.0225	9.62
		60.0	0.050	0.173	1.53	0.0225	9.44
		90.0	———{ *pending curation* }———
		120.0	———{ *pending curation* }———
		150.0	-0.087	0.100	1.50	0.0225	8.86
	0.4	30.0	0.173	0.200	1.55	0.0225	10.00
		60.0	0.100	0.346	2.32	0.0225	9.71
		90.0	-0.000	0.400	4.03	0.0225	9.09
		120.0	-0.100	0.346	2.97	0.0225	8.71
		150.0	-0.173	0.200	2.36	0.0225	8.45
	0.6	30.0	0.260	0.300	2.71	0.0225	7.36
		60.0	0.150	0.520	1.26	0.0225	9.81
		90.1	0.000	0.600	1.60	0.0225	9.10
		120.0	———{ *pending curation* }———
		150.0	-0.260	0.300	1.66	0.0225	8.04
	0.8	30.0	0.346	0.400	2.15	0.0225	10.95
		60.0	0.200	0.693	2.04	0.0225	10.28
		90.1	-0.000	0.800	2.07	0.0225	9.10
		120.1	-0.200	0.692	1.73	0.0225	8.29
		150.0	-0.347	0.399	1.15	0.0225	7.60
2	0.2	30.0	0.115	0.100	1.54	0.0225	10.59
		60.0	0.067	0.173	1.29	0.0225	10.26
		90.0	0.000	0.200	1.39	0.0225	9.85
		120.0	-0.067	0.173	1.79	0.0225	9.41
		150.0	-0.116	0.100	1.34	0.0225	9.12
	0.4	30.0	0.231	0.200	1.40	0.0225	11.14
		60.1	0.133	0.347	1.83	0.0225	10.78
		90.1	-0.001	0.400	1.30	0.0225	9.92
		120.1	-0.134	0.346	1.93	0.0225	8.95
		150.1	-0.231	0.200	1.50	0.0225	8.50
	0.6	30.1	0.346	0.301	1.42	0.0225	11.99
		60.1	0.199	0.502	1.68	0.0225	10.85
		90.1	-0.001	0.600	1.46	0.0225	10.23
		120.1	-0.201	0.519	1.71	0.0225	8.48
		150.1	-0.347	0.299	1.23	0.0225	8.36
	0.8	30.1	0.461	0.401	2.76	0.0225	12.45
		60.1	0.265	0.694	1.99	0.0225	11.43
		90.2	-0.002	0.800	2.80	0.0225	10.09
		120.2	-0.268	0.691	1.68	0.0225	8.31
		150.1	-0.462	0.398	1.37	0.0225	7.23

TABLE 4.2   *Initial parameters of the simulations presented in the catalogue. Eccentricity e is calculated over the region $200 - 1000M$ using the method described below. Orbital frequency $M\omega_{orb}$ is calculated from the dynamics after the passage of the junk radiation, with the number of orbits N_{orb} is defined between the time at which $M\omega_{orb}$ is calculated and peak in the $(2,2)$ mode of the Weyl scalar Ψ_4. Continued in* TABLE 4.3

q	χ_2	$\theta\,^\circ$	χ_{eff}	χ_p	e $(\times 10^{-3})$	$M\omega_{orb}$	N_{orb}
4	0.2	30.0	\multicolumn: ———{ *pending curation* }———				
		60.0	0.080	0.173	2.37	0.0242	11.09
		90.1	0.000	0.200	1.35	0.0255	9.36
		120.0	-0.080	0.173	1.22	0.0235	10.67
		150.0	-0.139	0.100	1.00	0.0233	10.37
	0.4	30.1	0.277	0.201	1.40	0.0252	11.96
		60.2	0.159	0.347	1.24	0.0246	11.20
		90.2	-0.001	0.400	1.44	0.0225	12.35
		120.2	-0.161	0.346	1.44	0.0225	10.96
		150.1	-0.277	0.199	2.03	0.0225	9.81
	0.6	30.1	0.415	0.301	1.68	0.0258	12.46
		60.2	0.238	0.521	1.39	0.0250	12.05
		90.3	-0.002	0.600	1.89	0.0239	10.71
		120.3	-0.242	0.518	1.03	0.0229	10.13
		150.0	\multicolumn: ———{ *pending curation* }———				
	0.8	30.2	0.553	0.402	1.35	0.0226	13.25
		60.4	0.317	0.695	0.75	0.0254	12.12
		90.4	-0.004	0.800	1.76	0.0225	12.33
		120.4	-0.324	0.690	2.01	0.0225	9.99
		150.2	-0.555	0.397	1.16	0.0225	8.15
8	0.2	30.0	0.154	0.100	1.30	0.0280	12.44
		60.0	0.089	0.173	1.00	0.0276	12.15
		90.0	0.000	0.200	0.89	0.0271	11.67
		120.0	-0.089	0.173	0.70	0.0265	11.37
		150.0	-0.154	0.100	1.59	0.0261	11.03
	0.4	30.2	0.307	0.201	1.25	0.0291	13.21
		60.3	0.176	0.347	2.11	0.0282	12.27
		90.3	-0.002	0.400	1.08	0.0272	11.79
		120.3	-0.179	0.345	1.34	0.0256	11.14
		150.1	-0.308	0.199	1.49	0.0249	10.66
	0.6	30.2	0.461	0.302	1.63	0.0302	14.29
		60.0	\multicolumn: ———{ *pending curation* }———				
		90.5	-0.004	0.600	2.08	0.0271	12.00
		120.4	-0.270	0.518	1.15	0.0255	10.84
		150.2	-0.463	0.298	0.82	0.0244	8.60
	0.8	30.3	0.614	0.404	1.01	0.0315	15.87
		60.5	0.350	0.697	1.36	0.0294	14.22
		90.6	-0.008	0.800	2.14	0.0271	11.92
		120.5	-0.361	0.689	2.75	0.0247	11.11
		150.3	-0.618	0.397	0.92	0.0233	10.18

TABLE 4.3   *Continuation of* TABLE 4.2

and the ratio of the root mean square norms of these errors in our two simulations in the limit of infinite resolution (that is, a completely continuous grid, no longer discrete) defines the *convergence factor* $c(t)$,

$$
\begin{aligned}
\lim_{d \to 0} c(t) &= \frac{||u(t) - u_1(t)||}{||u(t) - u_2(t)||} \ , \\
&= \frac{||k(t)d_1^n + \mathcal{O}(d)1^{n+1}||}{||k(t)d_2^n + \mathcal{O}(d)1^{n+1}||} \ , \\
&= \left(\frac{d_1}{d_2}\right)^n \ ,
\end{aligned}
\tag{4.6}
$$

where n is the order of the finite-difference scheme being used. But without access to that true solution we cannot compute the errors. Instead, we make use of three resolutions and examine the convergence behaviour of their relative differences,

$$
\begin{aligned}
\lim_{d \to 0} c(t) &= \frac{||u_1(t) - u_2(t)||}{||u_2(t) - u_3(t)||} \ , \\
&= \frac{d_1^n - d_2^n}{d_2^n - d_3^n} \ .
\end{aligned}
\tag{4.7}
$$

In practice we can plot the relative errors together, apply an appropriate rescaling and compare them by eye to estimate convergence. The infinite-resolution limit of an order n approximation scheme leaves each of our approximate solutions proportional to the n^{th} order error function, so for n^{th} order convergence

$$
u_1(t) - u_2(t) = \frac{d_1^n - d_2^n}{d_2^n - d_3^n} \left(u_2(t) - u_3(t)\right) \ ,
\tag{4.8}
$$

and the two plotted lines should lie on top of one another.

Accordingly the four simulations listed above are run at three resolutions, $d_{low} = 0.0125M$, $d_{med} = 0.0104M$, and $d_{high} = 0.00833M$, and in all cases the gravitational radiation is extracted at a selection of radii on the same refinement level: $r_{ext} = \{50, 60, 70, 80, 90\} M$. Note that the labels *low*, *medium*, and *high* are only in reference to their grid spacings relative to one another, not intended as a general statement on their accuracy in a wider numerical relativity context.

[85] finds fourth order convergence with respect to resolution and first order convergence with respect to the extraction radius. Using this an estimate is formed of the

resulting truncation error in each case. To calculate this the waveforms are rotated into the coprecessing frame (see §3.4), taking advantage of the freedom of choice in third Euler angle to align the phases at merger, and the amplitude, phase, and Euler angles α and β are extrapolated to an infinitely fine grid spacing using 4·6

$$u_{d\to0}(t) = \lim_{d\to0} \frac{c(t)u_2(t) - u_1(t)}{c(t) - 1} \,.$$

4·9

The percentage errors in amplitude, phase, and precession angles for the $(8, 0.8, 150°)$ case are shown in TABLE 4.4 and are representative of those in all configurations: of the order of a few percent in amplitude and phase, and around half a percent in precession angles. Thus the precessing waveform models that rely on these waveforms for calibration are limited to an accuracy of ±0.5%.

Source	$\Delta\phi$ (%)	ΔA (%)	$\Delta\alpha$ (%)	$\Delta\beta$ (%)
d	0.1	0.5	0.3	0.5
r_{ext}	1	8	0.03	0.5
Total	1	8	0.3	0.5

TABLE 4.4   *A representative example of the percentage errors in amplitude, phase, and precession angles. Shown here are the results for the* $(8, 0.8, 150°)$ *simulation.*

Primarily though these errors in A, α, β are a result of the dephasing. Certainly these are useful measures if we are comparing accuracy in simulations all generated by the same code as the dephasing error will be of the same magnitude and so contribute an overall shift, but beyond this they offer no particularly insightful commentary on accuracy in the context of waveform modelling, detector calibration, or any wider gravitational wave context. It is far more meaningful for the gravitational wave astronomy community to quantify accuracy in terms of the match, and accordingly in [82] it is found that the poorest match between a medium resolution waveform and an infinitely well resolved one is 0.9994, while that between a waveform extracted at $r_{ext} = 90M$ and at $r_{ext} \to \infty$ is 0.9986.

4.4. *Discussion*

We have presented here the first catalogue of black hole binary simulations generated using BAM by the Gravity Exploration Institute at Cardiff. Of the 80 simulations in the catalogue, four representative cases have been shown to demonstrate fourth-order convergence in resolution and first-order convergence in extraction radius. The error in precession angles has been found to be $\mathcal{O}$ (0.1%), and the minimum match between a waveform and its extrapolated forms exceeds 0.995. Any precessing model based on the simulations presented in this catalogue will therefore be limited by these uncertainties, assuming any other catalogues used possess no greater uncertainties.

There are a number of directions the next set of BAM simulations could take at Cardiff. One option would be to simply fill in more points within the subset of the parameter space covered by this catalogue: we would then have mass ratios $q = \{1, 2, 3, 4, 5, 6, 7, 8\}$ perhaps, and spins $\chi_2 = \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$, and so on. But with a variety of waveform models already developed that span this range, and even a model specifically tuned to the waveforms in this catalogue [61], all demonstrating a high degree of consistency in their reported PE results irrespective of whether they include precession effects or higher modes [86, 87], this is deemed to be excessive. Alternatively we may choose to extend this systematic coverage presented here, perhaps for example dedicating all our computing resources to the production of a few carefully placed simulations at very high mass ratios. But given the expense these would require[2] and the limited gain they would offer given current detectors' poor sensitivity to low frequencies (and therefore high mass ratios), this is dismissed. Instead an omission in the simulations here is addressed: the lack of spin on the secondary black hole. This will allow us to quantify the impact of spin-spin coupling on the morphology of the waveform. Although expected to be small, as we approach the era of third-generation gravitational wave detectors there will be increasing demand for waveform models that incorporate these subdominant effects in order to exploit the high-precision observations they will allow for.

With the publication of the BAM catalogue, the wider gravitational wave astronomy community now has access to simulated waveforms generated by four different numerical relativity codes; codes that each employ a variety of different techniques in their approach to evolving the black hole binary spacetime. The SpEC simulations in particu-

---

[2]See CHAPTER 6 for a more detailed discussion.

lar employ a radically different numerical setup to that outlined in PART I and used by
BAM, LazEv, and Maya. Rather than evolving moving puncture initial data through finite
differencing methods, SpEC uses a numerical grid that extends no further than the appar-
ent horizon of its black holes to avoid the singularity and solves the evolution equations
with pseudospectral methods.

It is curious then that very few consistency checks between the results of these codes
have been performed, particularly in light of their fundamental role they play in gravita-
tional wave astronomy. To what degree can, say, the $q = 8, a = 0.8, \theta = 150°$ binary
we've evolved for the BAM catalogue be said to agree with the same configuration evolved
using another code? Do the waveforms they generate agree sufficiently for the purposes
of current and future gravitational wave astronomy? It is precisely this question that we
aim to address in the following chapter.

A Cross-code Consistency Analysis

5

The first direct detection of gravitational waves by the LIGO-Virgo Collaboration in 2015 represented an important validation of general relativity and the culmination of decades of theoretical, experimental, and computational work. But six years on with 50 published detections the focus of the now LIGO-Virgo-KAGRA Collaboration has shifted to *gravitational wave science as astronomy*: the extraction of precision astrophysical results from these signals, particularly as we approach the third generation of detectors. It is therefore increasingly important that the tools used in the analysis of detected signals remain accurate in the face of the unprecedented signal power they will offer. Currently a detected gravitational wave is compared to a bank of model waveforms to see which one best resembles the signal, with these models informed by both analytic approximations and simulations that numerically solve the Einstein field equations. But if waveform models based upon numerical simulations are to be used in the analysis of the sources of detected gravitational wave signals, it is crucial that the accuracy of these simulations is thoroughly quantified and understood. Any biases in these simulations have the potential to propagate out through the waveform models they inform into the science output of the LVK, and with ever-increasing signal strengths anticipated in future detectors they will have an increasingly significant impact on gravitational wave astronomy. Typically the numerical waveforms used within the LVK are published together with some estimate of their uncertainties, but in this chapter we aim to further contextualise these uncertainties through the cross-validation of simulations generated by different numerical relativity codes, each employing different numerical methods, initial data, gauge choices, and indeed different formulations of the Einstein equations. This is by no means the

first such study, but is the first to investigate *precessing* black hole binaries and indeed the first to analyse spinning black holes in terms directly relevant to gravitational wave astronomy rather than just numerical relativity. This topic was explored in [88] from a purely numerical relativity perspective, followed later by the first investigation of the bearing such comparisons may have on gravitational wave detections in [89], and more recently a targeted comparison of numerical waveforms that mimic the GW150914 gravitational wave signal was conducted in [90]. This study advances these efforts to the modern age of gravitational wave astronomy in which previously subdominant features in the waveform play an increasingly significant role in accurate parameter estimation efforts.

In the third observing run of the LVK network we have seen the first strong evidence for mass ratios beyond unity and for the impact of higher harmonics in the signal, and the first hints of non-aligned spins [91]. In this chapter we aim to quantify the differences between precessing, unequal-mass simulations generated by three numerical relativity codes, and to contextualise the significance of these differences for gravitational wave astronomy. We explore the impact of differing gauge choices on precession dynamics, and analyse how this in turn affects the gravitational wave output in terms of the minimum signal-to-noise ratio (SNR) at which these differences would manifest with respect to the LIGO detectors. We examine both the gravitational waveform generated by the simulated binaries and the dynamics of the binaries themselves in our analysis, focusing on the final 12 orbits before merger from approximately $2000M$ before merger to around $200M$ after, where M is the total mass of the binary in geometric units, which translates to approximately $0.01(M/M_\odot)$ seconds. The simulations used in this analysis have been generated by the BAM, SpEC, and LazEv codes. BAM [44] and LazEv [47] are finite-differencing adaptive mesh refinement codes employing the moving puncture scheme in the BSSN formulation of Einstein's equations, while SpEC [92] uses pseudospectral methods to solve the generalised harmonic system for excision initial data.

This work builds on that of Hannam et al. [89] which made use of two of the codes we also use here—BAM and SpEC—to study the consistency of non-spinning equal-mass simulations. It was found that both the phase and amplitude disagreements are within the uncertainty estimates, resulting in a match (in the (2,2) mode) for all codes above 0.999: indistinguishable at signal-to-noise ratios below ≈ 14. More recently Lovelace et al. in [90] compared simulations of an aligned-spin near-equal-mass binary representing

the probable source of GW150914 generated by the SpEC and LazEv codes. Matches here are also seen to exceed 0.999 across the frequency range, with differences in the gravitational wave phase found not to exceed 0.01 rad over the final six seconds before merger.

However astrophysical black hole binaries are generically expected to be spinning, with no reason to suppose *a priori* that those spins are aligned with the orbital plane. In such cases the spin-orbit and spin-spin couplings lead to precession of the orbital plane, in which the orbital angular momentum vector rotates about the total angular momentum vector. Compared to non-precessing configurations this brings added complications to efforts to numerically simulate these binaries as the rapidly changing phase near merger translates to significant short-timescale rotation of the orbital plane that must be captured accurately. Additional complications plague the behaviour of the resulting gravitational waveform, covered in detail in [56, 63, 93]. Briefly, the action of subdominant modes in the gravitational wave signal becomes significant as the observed power is shunted into them from the dominant (2,2) mode—indeed, there are configurations where the otherwise-subdominant harmonics dominate [63]. There has been to date no consistency check of the extraction of these higher modes between numerical relativity codes, nor has there been any exploration of the relative phasing between precessing waveforms or of the dynamics of the orbital plane.

While certainly of interest to the numerical relativity community, up until now this has been largely irrelevant to gravitational wave astronomy. During the first two observing runs of the LIGO-Virgo detector network there was no evidence of precession in the detected signals [80]. The third observing run has resulted in the strongest hint yet of precession [91] and the detection of the first binary where the measurement of the mass ratio was clearly bounded away from $q = 1$ [81, 87], while recent developments in waveform modelling have, using precessing numerical simulations, produced models that include the effects of precession [60, 94] that were used in the analysis of all O3 events. With the prospect of ever more detections of such binaries on the horizon it is crucial that the consistency of the precessing NR simulations from which they are built is quantified.

To illustrate this, TABLE 5.1 details the simulations used in the construction of some of the models within the EOB and phenomenological families. We can see that EOB models have for the most part been developed using, and calibrated against, excision simu-

lations, while the `IMRPhenom` family include in their construction simulations produced by moving-puncture codes. Now given that these codes differ significantly in numerical methods, initial data, gauge choices, and evolution schemes one can imagine it possible that inconsistencies in the code would propagate out through the numerical waveform to the waveform models, and so out to parameter estimation efforts and the science output of the LVK. It has been demonstrated in [89] that the `BAM` and `SpEC` codes, as they existed a decade ago, are sufficiently consistent in the dominant (2,2) mode for the current generation of ground-based detectors for non-spinning black holes with equal mass components. This has been extended to spinning black holes in [90], focusing in detail on one example of how simulation accuracy may affect the analysis of a detected signal. But there has as yet been no investigation of their consistency in light of the more recent proliferation of unequal mass, precessing simulations. To accurately represent the waveform in such cases there is now a need for much tighter consistency in the phase evolution across different codes, and the accurate extraction of the subdominant modes—in short, there are a great many more areas where differences between codes may show up than in the equal-mass aligned-spin comparisons performed in the past.

| | SpEC | Moving puncture | | |
		BAM	Llama	CCATIE
PhenomA [95]	-	3	-	1
PhenomB [96]	-	23	2	8
PhenomC [97]	1	23	2	16
PhenomD [70]	174	18	-	-
PhenomP [60]	174	18	-	-
PhenomHM [63]	174	18	-	-
SEOBNRv2 [98]	38	-	-	-
SEOBNRv3 [71]	70	-	-	-
SEOBNRv4 [72]	140	1	-	-
SEOBNRv4HM [73]	140	-	-	-
SEOBNRv4PHM [74]	1523	-	-	-

TABLE 5.1　*Simulations used in the construction of various EOB and phenomenological models, not including those used for verification purposes.* `PhenomP` *and* `PhenomHM` *are both based on* `PhenomD`, *hence the identical numbers of simulations used.*

The chapter is organised as follows. We first outline in §5.1 some details of three different codes whose simulations are currently used in the development of waveform models used in searches and parameter estimation, focusing on numerical techniques, initial data, evolution schemes, and gravitational wave extraction. §5.2 introduces the specific physical scenario that will be compared across these codes as well as the preprocessing that has been done to get them sufficiently aligned with one another, together with a brief description of the various challenges involved in such an alignment, and in §5.3 and §5.4 we present the results of this comparison. This covers both the dynamical evolution of the black holes and their associated gravitational waves, together with calculations of the match between the waveforms in §5.5. Finally §5.6 is given over to discussion of these results in the context of upcoming advances in gravitational wave astronomy.

5.1. *Numerics*

As we have outlined in CHAPTER 1, in order to evolve dynamical spacetimes numerical relativity aims to recast the Einstein field equations in the form of an initial value problem, comprising a fully described initial state of the spacetime together with a set of partial differential equations whose solutions represent the state at some future time. We compare three numerical relativity codes—BAM, LazEv, and SpEC—that approach this problem in a number of different ways, summarised in TABLE 5.2. The BAM simulations were run by the author, the LazEv simulations by Jim Healy, and the SpEC simulations through the combined efforts of Katerina Chatziioannou, Harald Pfeiffer, and Geoffrey Lovelace.

To ensure the analysis can freely be applied to any simulation that may be produced for the LVK, and to standardise many of the convention choices involved in the presentation of the data, we choose to work exclusively with simulations that have been curated according to the LVC NR Injection Infrastructure (NRII), a framework outlined in [99] that standardises many of these choices and sets out the format in which all NR simulation data is to be stored: a single HDF5 file containing time domain datasets of specific, unambiguously defined quantities. For this study each simulation is provided following the 'Format 2' requirements, containing phase and amplitude time series of each mode of the gravitational wave strain $h_{\ell m}$ as well as time-domain dynamics data: the black holes' positions, spins, orbital frequency and so on.

	Formulation	Initial data	Topology	Gauge
BAM	BSSN	Bowen-York	Wormhole	1+log $\widetilde{\Gamma}$-driver
LazEv	BSSN	Bowen-York	Wormhole	1+log $\widetilde{\Gamma}$-driver
SpEC	Generalised harmonic	Conformal thin-sandwich	Excised	Damped harmonic

	Time-stepping	Mesh	Discretisation scheme	Extraction radius
BAM	Runge-Kutta (4$^{\text{th}}$ order)	Berger-Oliger nested Cartesian	Finite difference	50, 60, 70, 80, 90, 100, 110, 120
LazEv	Runge-Kutta (4$^{\text{th}}$ order)	Berger-Oliger nested Cartesian	Finite difference	∞
SpEC	Adaptive Runge-Kutta (4$^{\text{th}}$/5$^{\text{th}}$ order)	Nested spherical-cubic-spherical	Pseudo-spectral	∞

TABLE 5.2   *Summary of the three codes featured in this study, based partly on Table I of [90]. An extraction radius given as ∞ indicates extrapolation to infinity.*

Even before any discussion of details specific to numerical relativity and the various choices available to us there, there is a much more basic level at which these codes differ, namely the manner in which they numerically solve the sets of differential equations describing the evolution of the spacetime. The codes we work with here fall into two broad camps. BAM and LazEv make use of finite-differencing (sixth and fourth order respectively) as described in §2.5 in which continuous derivatives are replaced with discrete differences and the evolution equations are represented by a system of linear equations, while SpEC employs pseudospectral methods to solve the evolution equations. In contrast to finite-differencing this approach uses a linear combination of basis functions to represent the solution to the evolution equations. The coefficients of the basis functions are time-dependent, and are themselves solved as simply a system of ordinary differential equations. In practical terms this approach offers significantly improved accuracy and efficiency displaying exponential convergence, with the downside of extreme sensitivity to instability [100]. Time-stepping is performed by BAM and LazEv through fourth-order Runge-Kutta integration with Berger-Oliger refinement of the discretisation scheme for the finer levels, and in SpEC through a fifth-order Dormand-Prince integrator. In the

latter case the size of the time steps is chosen so as to constrain time-stepping errors to within a given tolerance by an adaptive proportional-integral control system [79].

Already we see how simple choices even at this level have the potential to produce markedly different results: one can imagine, for example, that the poorer convergence properties of the finite-difference codes may lead to a far more inaccurate simulation by the time we reach merger, or perhaps that some unnoticed instability could afflict the more sensitive spectral code. Beyond this initial branch point are a number of further forks in the road that each have the potential to lead to significantly different simulations.

For one thing, although all three codes make use of adaptive mesh refinement, the coordinate structure of the numerical domain differs significantly. The finite-difference codes both follow the structure outlined in §2.5, employing a series of nested Cartesian grids of increasingly fine spacing and extent, one set of grids centred on each puncture, where upon reaching a refinement level where the two black holes' grids would overlap they are replaced by a single grid that is of sufficient size to contain the two. SpEC features a domain consisting similarly of nested grids, again with finer levels centred on each black hole and coarser levels containing both, though with a rather more complicated structure. Around each black hole is a set of concentric spherical shells, progressively deforming into a cube with each additional layer. As we continue further outwards these two cubes are themselves surrounded by cuboids in order to fill up a larger cube, which in turn is surrounded by concentric layers this time steadily deforming into spherical shells (see [101] for further details). In fact during the evolution a dual-frame method is used in which the basis of the various evolved fields is constructed on a static asymptotically Cartesian grid while the evolution equations are solved on a co-moving frame that tracks the motion of the horizons.

Both BAM and LazEv evolve Bowen-York initial data. In contrast, SpEC employs conformal thin-sandwich initial data extending only as far as the horizon [79, 101], with the region interior to the apparent horizon entirely excised from the computational domain. The conformal thin-sandwich data utilises a modification of the York-Lichnerowicz conformal decomposition to solve the constraint equations, whereby the conformal metric on two surrounding slices is provided as free data along with the conformal metric on a given slice, the trace of the extrinsic curvature, and the energy and momentum density [10]. Whereas Bowen-York initial data allows the spins of the black holes to be specified directly, thin-sandwich data must be constructed with appropriate boundary

conditions on the shift vector at the outer boundary and each apparent horizon such that the desired configuration is produced [102–105].

For BAM and LazEv the desired spin orientations of the black holes are achieved through the iterative algorithm presented in [106] and summarised in CHAPTER 4. Briefly, the orientation is provided initially at large separation and the system is evolved via the EOB equations of motion up to the desired reference orbital frequency at which the simulation will begin. Of course by doing this we end up with a slightly different spin configuration at the start of the simulation than we were aiming for because the spin will have evolved in the intervening time, so the difference between the actual and desired spin directions is noted and the initial orientation is rotated by that amount. This is repeated until initial parameters are obtained that produce the desired configuration at the starting frequency to a tolerance of 1%.

One aspect common to all three codes is the need for some eccentricity reduction procedure on the initial data so as to generate a sufficiently quasi-circular binary. Again, this has been discussed in further detail in CHAPTER 4 and is achieved by iteratively applying small adjustments to the initial parameters. The two moving-puncture codes follow the procedure outlined in [84], fitting the puncture separation during early inspiral using a quadratic function and defining the eccentricity as the maximum difference between the fit and the data in this range.[1] For BAM the fit is made over the range 200–700M, after which small perturbations are introduced to the magnitude of the black holes' linear momenta until the eccentricity is below 0.002 over the first 1000M of simulation time. If this cannot be achieved through perturbation of the linear momenta alone, the radial momenta are decreased.

In addition to these broader differences in approach to initial data construction in general, the particular initial data used in the simulations produced for this study differ, and are crafted in order to produce a particular orientation and spin configuration defined in §5.2 some time into the simulation. This is in order to avoid the initial junk radiation in each simulation, handled in the final curated data product by simply cropping the first few hundred M of the simulation time.

The differing impact of junk radiation across these codes is one potential source of errors introduced by the radiative content of the simulation. Even if it were eliminated

---

[1]The full method described in [84] requires the merger time, and is used to estimate the eccentricity of the final simulation.

entirely there remain differences in the precise way in which the gravitational wave signal is extracted from the simulation. In all three codes the radiation content of a simulation is computed as described in §2.4 at a number of different radii with an associated error resulting from this extraction at finite distance, rather than at spatial infinity where it is defined, but beyond this all three codes differ. No further treatment of this error is performed by BAM, while SpEC and LazEv on the other hand mitigate this issue by proceeding to extrapolate an approximation of the waveform at infinity by first aligning the waveform extracted at each radius in time, and treating the error as a polynomial in $\frac{1}{r_{ext}^n}$. For SpEC the extraction radii used in this extrapolation vary from one simulation to another, but typically begin at $100M$ and extend to the outer boundary in 24 steps uniform in $\frac{1}{r_{ext}}$ [79]. LazEv uses ten radii spaced between $75M$ and $190M$ in the same way [31, 107].

5.2. *Setup & Alignment*

As the black holes orbit, the loss of energy through gravitational radiation both decreases their separation, and circularises the orbit [108]. As the black holes orbit closer, their angular velocity increases, which in turn increases the rate of energy loss, and the inspiral proceeds ever faster to merger. The highly perturbed remnant rapidly relaxes to a Kerr black hole through further gravitational wave emission, where the frequency and exponential decay rate of the ringdown signal can be related through perturbation theory to the mass and spin of the final black hole [109]. TABLE 5.3 lists the masses and spins of the five simulations we consider here.

It is often useful to decompose the gravitational wave signal of such an inspiralling binary into spherical harmonics, dominated by the ($\ell = 2, m = \pm 2$) modes in the non-precessing case. Taking the equal-mass non-spinning case as an example, the frequency (and amplitude) of these modes increase during inspiral from $0.38(M_\odot/M)$ kHz one hundred orbits before merger to $1.95(M_\odot/M)$ kHz in the final 5 orbits [89], before finally reaching their maxima around merger. After merger the amplitude of the signal decays exponentially in a process known as *ringdown*, where the perturbed remnant transitions to a Kerr black hole. The LVK operates gravitational wave detectors sensitive in the frequency band 10 Hz to 10 kHz [110] with the final six orbits (with merger) in-band for binary masses around $50M_\odot$. With this in mind we choose to investigate here simula-

Simulation	$m_{1,i}$ $(\pm 5 \times 10^{-7})$	$\chi_{1,i}$ $(\pm 5 \times 10^{-7})$	$m_{2,i}$ $(\pm 5 \times 10^{-7})$	$\chi_{2,i}$ $(\pm 5 \times 10^{-7})$
BAM	0.749917	0.750266	0.250037	2.44×10^{-5}
BAMHR	0.749791	0.750578	0.250044	1.27×10^{-5}
LazEv	0.749718	0.750632	0.250004	0.17×10^{-5}
SpEC	0.750153	0.749746	0.249973	9.21×10^{-5}
SpECHR	0.749918	0.748581	0.250005	7.28×10^{-5}

Simulation		M_f $(\pm 5 \times 10^{-7})$	χ_f $(\pm 5 \times 10^{-7})$
BAM		0.968062	0.488133
BAMHR		0.968303	0.487038
LazEv		0.966913	0.486429
SpEC		0.968673	0.489844
SpECHR		0.968648	0.489426

TABLE 5.3 *Initial and final masses and spins of the five simulations considered in this chapter.*

tions of the final twelve orbits of a quasi-circular black hole binary. We focus on a single configuration: a mass ratio $q = 3$ binary in which only the larger black hole is spinning, with dimensionless spin magnitude $\chi_1 = 0.75$ entirely in the orbital plane, aligned about an orbital frequency of $0.02/M$. The binary completes approximately half a precession cycle before merger.

Having introduced the coprecessing frame in §3.4, here we will need to be more meticulous in our choice of terminology to distinguish between frames defined with respect to the orbital dynamics and those defined with respect to the waveform, and so throughout will refer to the frame in which the Newtonian orbital angular momentum vector remains static as the $\hat{\vec{L}}_N$-*aligned frame*, and the frame which maximises at all times the amplitude of the (2,2) mode as the *quadrupole-aligned frame*, or the LA and QA frames for brevity. Now at first it would seem that these would be identical, but even upon their introduction in [57] this had been shown not to be the case. At the time it was presumed, with support from comparisons against PN estimates of $\vec{L}$, that this difference was a result of using the Newtonian quantity and that the 'full' angular momentum would align with the QA direction. But later developments [111, 112] have shown that the quadrupole-aligned direction does not track the orbital plane direction but depends both on the particular radiation frame used and whether it is calculated from the Weyl scalar Ψ_4, the Bondi news $\dot{\Psi}_4$, or the strain $\ddot{\Psi}_4$ [51]. In this work however we use the strain throughout, and

since we are making the same choice for all simulations we have a consistent definition of a quadrupole-aligned frame in which to make comparisons. We employ the precession angles of §3.4, (α, β, γ), to define the rotation into either the LA or QA frame depending on context.

In comparing the gravitational wave signal, the relative alignment of the system can significantly affect the apparent accuracy of the amplitude and phase evolution, as can the choice to parameterise these by time or by frequency [89]. If we were simply comparing identical simulations run on the same code differing only in resolution either of these would be a suitable choice and the alignment would be relatively trivial, but in comparing *between* codes, with different initial data and gauge conditions and conventions, poor alignment of waveforms can amplify or suppress true differences between the output of these codes. Naïvely we may choose merger as the event about which to align our simulations, presuming it to be well-defined and cleanly identifiable in the waveform as the moment that the amplitude of the (2,2) mode in the QA frame is maximised. We would then apply a time shift that sets this event at $t = 0$ and proceed to examine, say, the difference between the phases of our waveforms. Unfortunately merger time is not well-defined in the waveform [113], and so any synchronisation of two waveforms based on it—and the resulting differences observed—are highly sensitive to the definition used. If we were considering only the gravitational wave signal, some suitably well-defined reference point could be chosen and it would be possible to determine if the waveforms agree within their uncertainty estimates; but we both lack such robust waveform error estimates, and more fundamentally this might well show that the waveforms agree but the dynamics could display clear differences. These would then appear in the waveform data at higher accuracies so it is important to investigate the dynamics quantities as well. In the past a successful solution to this has been to compute the additional time shift that minimises the phase disagreement, but in this work we are looking to study these disagreements rather than suppress them. In short, there is a great deal of freedom in how we align the waveforms, and some alignments may act to reduce the magnitude of any differences that may exist between them making them harder to identify, while others will make these differences clearer. For this reason we aim to align the waveforms in the most well-defined way available to us that will also highlight differences that may exist.

To circumvent these difficulties [89] chooses to work in the frequency domain, parameterising quantities by the frequency of the (2,2) mode of Ψ_4 rather than by time,

thereby avoiding the problematic sensitivity to time shifts. Unfortunately this is not an option in the current study as per the NRII we only have access to the gravitational wave strain h, not Ψ_4, with the two related by 3·1. This is more of a problem than it may at first seem. There is in principle no reason we can't parameterise by the strain frequency, but in practice the double integration over the whole evolution required to generate h from Ψ_4 amplifies any noise in Ψ_4; particularly problematic for higher modes as they require greater resolution due to their higher frequencies. And given that we are working with a binary designed to display particularly strong precession, where there is considerable power in these modes, these amplified inaccuracies represent a significant problem. With the strain phase already compromised in this fashion, taking its derivative to obtain the frequency exacerbates the problem even further. One option here is to use the orbital frequency in place of the gravitational wave frequency—that after all is significantly cleaner, and as the phase of h is related to orbital phase by

$$\phi_h \approx 2\phi_{orb} + \pi \, , \qquad\qquad 5 \cdot 1$$

the frequencies of the two are linked by a simple factor of two, providing us with what is essentially a clean and relatively accurate fit through the noisy data. But the relation 5·1 is only an approximation, one which only holds during early inspiral and so would only serve to introduce further error in our attempts to locate a precise event. There is the additional, though lesser, problem that for moving puncture simulations the black holes remain tracked even through merger and ringdown, and so an orbital frequency can be computed; but of course for a code based on excision initial data this is not possible. We would then be left only being able to compare the SpEC simulation during inspiral. With all this in mind the quantities used in our comparisons remain parameterised by time. While the issues discussed with this approach do indeed remain we anticipate that they present a lesser risk than attempting to work with quantities parameterised by frequency.

With the frequency domain unsuitable and merger deemed not suitably well-defined, we must find some other event around which we can synchronise the simulations. To this end we require that in each simulation at an orbital frequency $M\Omega_{ref}^d = 0.022$ the spin of the large black hole must point directly at the smaller black hole, which should lie on the negative and positive x-axes respectively with the orbital plane entirely in
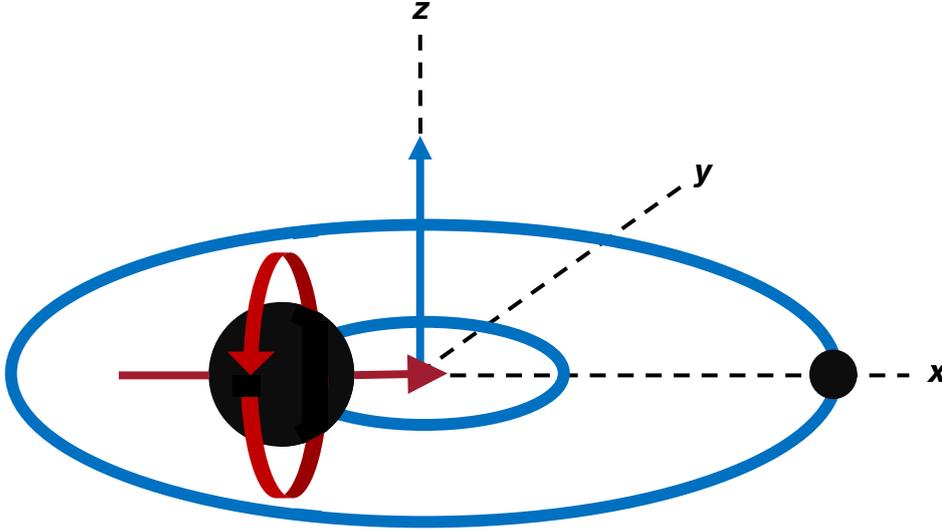
FIGURE 5.1   *Schematic outline of the configuration at the reference orbital frequency $0.022/M$. The blue tracks show the quasi-circular orbits of the black holes with the blue arrow as the normal to the orbital plane. The red arrows indicate spin.*

the xy-plane (FIG 5.1). Through this we define a relatively robust reference event about which the simulations can be aligned. This can of course never be perfectly immune to the relativity of simultaneity inherent to general relativity; there will always exist some coordinate system in which the two events—that is, hitting the right frequency and hitting the right spin configuration—are not simultaneous, but it is hoped that this event is such that the codes used here, all ultimately employing physically reasonable and relatively intuitive coordinates, will agree to within numerical precision.

The first task then is to identify this reference event in the simulations and rotate them accordingly. Identifying when this reference event occurs is not straightforward. There is no single 'simulation time' from which to select one entry where the event occurs. As we've seen in §5.2, the black hole dynamics data is calculated from their horizons while the gravitational wave data is extracted at some finite distance from the system and so will be delayed accordingly. The output of the simulations can then be split into two distinct parts: dynamics data and waveform data. This is problematic because there exists no trivial relation between the two—the dynamics and waveform time series, t_d and t_w, are not synchronised, so locating a time in one that corresponds to a specific time in the other is difficult. The solution may at first seem trivial: the difference between the two is just

the light travel time from the binary to the extraction radius. This would certainly be the case in flat spacetime, or would be sufficient in the weak-field limit, but around the binary an unambiguous relation between t_d and t_w does not exist.

Identifying the reference event in the dynamics alone is fraught with subtleties that need to be addressed. Conceptually it's straightforward enough: we simply locate when $M\Omega = 0.022$ in the orbital frequency data, which is one of the datasets required by the NRII defined in [99] as

$$\vec{\Omega}(t) = \hat{\vec{n}}(t) \times \frac{d\hat{\vec{n}}(t)}{dt} , \qquad\qquad 5\cdot2$$

where $\hat{\vec{n}}$ is the unit separation vector between the black holes. But being *quasi*circular these simulations have some small but non-zero eccentricity. Their orbits are slightly elliptical and so the black holes will be moving faster at the semi-minor axis than at the semi-major. This eccentricity therefore manifests itself as oscillations in the orbital frequency about the ideal circular orbital frequency (FIG 5.2), which can be approximated by the *orbit-averaged* orbital frequency.

Taking the BAM simulation as a representative example, if we try to identify when a particular orbital frequency occurs we see that a choice to work with either the instantaneous or the averaged quantity is—if we happen to land on a maximum in these oscillations—equivalent to a difference in Ω of as much as $10^{-4}/M$. While this may seem insignificant, it turns out that the configuration of our binary at the reference time t_{ref}^d is extremely sensitive to these eccentricity-induced oscillations. To see why, from the PN formalism we have that orbital frequency is given by [114]

$$\Omega(t) = \left(\frac{5}{256} \frac{(1+q)^2}{q} \right)^{\frac{3}{8}} \left(t_{merge} - t \right)^{-\frac{3}{8}} , \qquad\qquad 5\cdot3$$

and so in aligning two simulations in the manner described above, we are aiming to define the time t_{ref}^d such that $M\Omega(t = t^d ref) = 0.022$, given by

$$t_{ref}^d = t_{merge} - \frac{5}{256} \frac{(1+q)^2}{q} \Omega^{-\frac{8}{3}} \left(t_{ref}^d \right) . \qquad\qquad 5\cdot4$$

Repeating this for an eccentric orbit, where $\Omega \to \Omega + \Delta\Omega$, we find that the difference in time between an eccentric and a non-eccentric binary hitting the reference frequency goes as
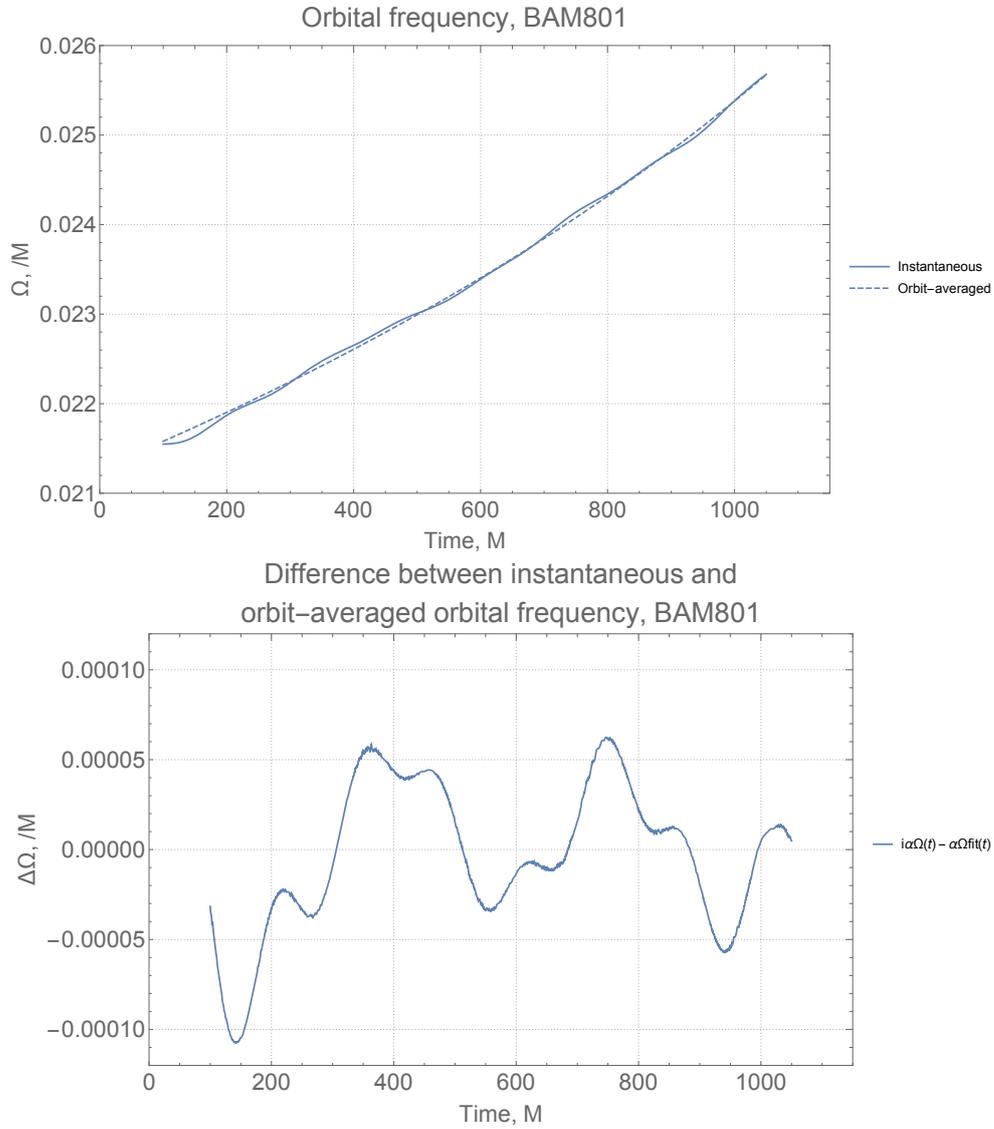
FIGURE 5.2 *Illustrating the impact of even small eccentricities on the orbital frequency. The top panel makes clear the oscillations induced by the eccentricity, with the instantaneous orbital frequency given by the dashed line and the orbit-averaged frequency given by the solid line. The difference between these two is taken in the bottom panel.*

$$\Delta t_{ref}^{d} \quad = \quad \frac{5}{256}\frac{(1+q)^2}{q}\left(\Omega^{-\frac{8}{3}}-(\Omega+\Delta\Omega)^{-\frac{8}{3}}\right), \tag{5.5}$$

$$\approx \quad \frac{5}{256}\frac{(1+q)^2}{q}\Omega^{-\frac{8}{3}}\left(1-\left(1-\frac{8}{3}\frac{\Delta\Omega}{\Omega}+...\right)\right), \tag{5.6}$$

$$\approx \quad \frac{5}{96}\frac{(1+q)^2}{q}\Omega^{-\frac{11}{3}}\Delta\Omega, \tag{5.7}$$

where higher-order $\Delta\Omega$ terms have been neglected. For our $q = 3$ system under consideration here, using the $\Delta\Omega = 10^{-4}$ of the BAM simulation seen above, at $M\Omega = 0.022$ this would produce a difference of $\Delta t_{ref}^{d} \approx 33M$. In order to determine the phase difference accumulated during this interval we recall that frequency is simply the time derivative of the phase, and so we can integrate 5.3 once more to obtain

$$\phi_{orb}(t) = -\frac{8}{5}\left(\frac{5}{256}\frac{(1+q)^2}{q}\right)^{\frac{3}{8}}\left(t_{merge}-t\right)^{\frac{5}{8}}+\phi_{ref}. \tag{5.8}$$

We are free to choose that $\phi_{ref} \equiv \phi_{orb}(t = t_{ref}^{d}) = 0$, and so

$$\phi_{orb}(t) = -\frac{8}{5}\left(\frac{5}{256}\frac{(1+q)^2}{q}\right)^{\frac{3}{8}}\left(t_{merge}-t\right)^{\frac{5}{8}}+\underbrace{\frac{8}{5}\left(\frac{5}{256}\frac{(1+q)^2}{q}\right)^{\frac{3}{8}}\left(t_{merge}-t_{ref}^{d}\right)^{\frac{5}{8}}}_{\phi_{ref}}. \tag{5.9}$$

The accumulated phase difference over the interval Δt_{ref}^{d} is then given by

$$\Delta\phi_{orb} = \phi_{orb}\left(t+\Delta t_{ref}^{d}\right)-\phi_{orb}(t), \tag{5.10}$$

$$= -\frac{8}{5}\left(\frac{5}{256}\frac{(1+q)^2}{q}\right)^{\frac{3}{8}}\left(\left(t_{merge}-t-\Delta t_{ref}^{d}\right)^{\frac{5}{8}}+\left(t_{merge}-t\right)^{\frac{5}{8}}\right), \tag{5.11}$$

$$\approx \left(\frac{5}{256}\frac{(1+q)^2}{q}\right)^{\frac{3}{8}}\left(t_{merge}-t\right)^{-\frac{3}{8}}\Delta t_{ref}^{d}, \tag{5.12}$$

$$\sim \Delta t_{ref}^{d}, \tag{5.13}$$

neglecting higher-order Δt_{ref}^{d} terms.

The spin vector will also have evolved over this interval—we see that for otherwise identical simulations, aligning one according to our earlier requirements at $0.022 / M$

in the *instantaneous* frequency and another at $0.022\ /M$ in the *orbit-averaged* frequency produces an angular separation of their spin vectors of 0.03 rad.

This is precisely the issue we are now confronted with. For the purposes of our alignment we work with the orbit-averaged frequency in order to minimise the impact of these eccentricity oscillations, but this choice, indeed either choice here, immediately presents us with a problem: the SpEC simulation has been crafted so as to achieve the correct spin orientation at $0.022\ /M$ in the orbit-averaged frequency, while BAM and LazEv instead have used the instantaneous frequency. Locating our reference frequency in the orbit-averaged data therefore sets spin offset at t_{ref}^{d} and we no longer have such clean alignment of the simulations.

This orbit-averaged orbital frequency is determined for each of the simulations we consider by a fit of the form of 5·3 over the range $t = \{200, 1200\}\ M$. With this in hand we apply a time shift to the dynamics data such that $t_{ref}^{d} = 0$. Then, taking the data at this time, we find and apply the rotation matrix that will put the normal to the orbital plane along the z-axis. A second rotation then puts the black holes on the x-axis with the smaller black hole on the positive axis, which is simply a two-dimensional rotation about the z-axis of $-\phi_{orb}(t = 0)$ to set the orbital phase of the binary to 0. Combining these rotation matrices we now have a single rotation that orients the binary as we would like, which we apply at all times to the original data.

At this stage we have the dynamics oriented such that when the orbital frequency is $0.022\ /M$, the black holes lie on the x-axis with their orbital plane entirely in the xy-plane. To find the corresponding reference time t_{ref}^{w} in the waveform data we apply a rotation to the waveform data that roughly corresponds to that which we have just applied to the dynamics and look for the moment when, based on the waveform data, this orientation is achieved. The first step is to rotate the waveform into the QA frame— which, recall, is approximately equivalent to putting the dynamics into the LA frame earlier—achieved by finding the set of rotations that at each moment maximise the amplitude of the (2,2) mode. This is performed using the nrutils suite developed by Lionel London. We then apply the second rotation that was applied to dynamics, that corresponded to a final phase shift. But some care must be taken here: the orbital phase ϕ_{orb} and the phase of the gravitational wave strain ϕ_{h} are not identical, but related by 5·1,

so the second rotation is now a shift of

$$\Delta \phi_h = -\left(2\phi_{orb}(t_{ref}^d) + \pi\right) .$$

We know that the QA direction is not in general the same as the LA direction, and so it almost certainly does not precisely correspond to the orbital plane being the xy-plane, but unlike the gauge-dependent black hole positions it is at least a well-defined direction. The invariant nature of Ψ_4 means that using this method would perfectly align the waveforms produced by each code with one another (assuming identical initial data and infinite precision), whereas using the dynamics may not due to the different gauge choices in each code. Given that the waveform data and dynamics data are now in approximately equivalent frames, with the black holes set on the x-axis at the reference time, we can narrow down our search for t_{ref}^w to just those times when the waveform data shows the black holes crossing the x-axis. As the orbital phase is $n\pi, \{n \in \mathbb{Z}, n \geq 0\}$ when the black holes cross the x-axis[2] we know from 5·1 that this translates to a gravitational wave phase of $(2n + 1)\pi$, so we search for every instance that this occurs.

In fact we can narrow this further as we know that the small black hole must be on the positive x-axis at the reference time. The other crossing of the x-axis, with the small black hole at negative x, will not do, and searching in the waveform phase for $(2n + 1)\pi$ does not discriminate between these two scenarios. This means we're really looking for an orbital phase of $2n\pi$, and so a gravitational wave phase of $(4n + 1)\pi$. We then look for all occurrences of this and build a list of potential times that could correspond to the reference event in the dynamics. Which of these is the one we want? This is potentially problematic: aside from the issues mentioned above regarding ambiguity in times, there is also the fact that the waveform data is cropped to remove the initial ripple of junk radiation at the start of the waveform. There is no way of telling, from the NRII data alone, how much time has been cropped, or if there has been any padding of the waveform at the start and end of the data. Thankfully though the NRII requires that the dynamics and waveform data be approximately synchronous, a requirement that also takes care of the time shift due to extracting the waveform far from the binary, so we can search for the time in the list that is closest to the dynamics reference time. As each orbit, and so

---

[2]This is true only for `BAM` simulations—for `SpEC` and `LazEv` this phase occurs when the black holes cross the y-axis instead, and as such appropriate modifications must be made throughout to account for this.

each entry in our list of candidate times, is separated by approximately 300 M, even if the simulation follows the NRII requirements only very roughly we are unlikely to have settled upon the wrong time.

Now that we have t_{ref}^w, we select the single rotation that was applied at that instant and apply it to the original data at all times, just as we did for the dynamics. Finally, a time shift is applied such that $t_{ref}^w = 0$. Our dynamics data and waveform data are now oriented identically and synchronised.

It is worth noting, however, that the relation 5·1 between orbital and strain phases only holds during early inspiral and to leading order in the (2,2) mode only, so there is the potential for some inaccuracy in t_{ref}^w. We can investigate this by working the other way round: identifying t_{ref}^w first, finding the rotations at that point, then applying that to the dynamics data and inferring the corresponding t_{ref}^d. If we set t_{ref}^w to the time found in the dynamics-first approach we should recover the same t_{ref}^d if 5·1 is sufficiently valid. We find that it differs by $\sim 30M$, corresponding to a phase difference of 0.03 rad in the orbit, of a similar order to the phase uncertainties due to the effects of eccentricity on the orbital frequency discussed above.

We now turn to the spin angular momentum to examine how well each simulation has approximates the configuration we are aiming for. FIG 5.3 shows the track of the spin vector across the sky throughout the entire evolution of the large black hole. Recall that in setting up these simulations it was intended that the spin on the large black hole lie directly along the separation of the two black holes at $t = 0$, which would correspond to the dashed lines in FIG 5.3 pointing towards $\frac{\pi}{2}$ rad—clearly not the case here. This is a result of the previous discussion regarding the precision to which the configuration of the binary is crafted given the ambiguity in orbital frequency introduced by non-zero eccentricities. We can convince ourselves though that this initial difference, this initial misalignment between the spins in these simulations, remains constant throughout the inspiral rather than contributing some growing error if we abandon the alignment in phase at $t = 0$ we have so far favoured and instead rotate the frames such that the spins agree at this moment. The results of this are also shown in FIG 5.3, and we find that the spins remain aligned to a tolerance of 0.05 rad between codes, and to within 0.005 rad between simulations run with the same code.

The synchronisation of the spins throughout the evolution of the binary is rather remarkable. Between the BAM and SpEC runs we have different gauges, different topology,
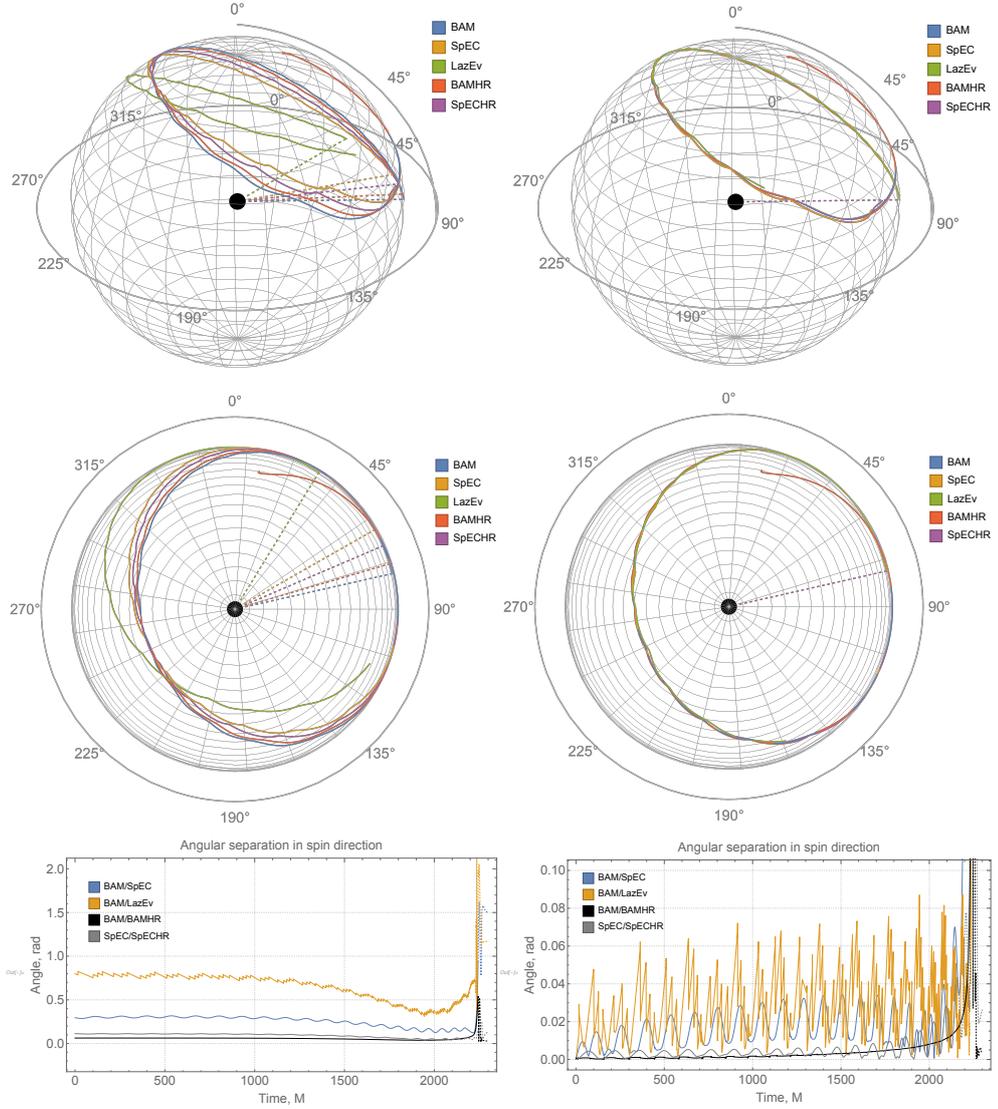
FIGURE 5.3   *The tracks of the spin angular momentum in each simulation, with phases aligned at t_{ref}^{d} in the left-hand column and spins aligned at t_{ref}^{d} in the right-hand column. The black sphere in the centre represents the spinning black hole while the dashed lines show the value at t_{ref}^{d}.* TOP: *3D representation of spin tracks on the sky.* CENTRE: *Top-down view of the same.* BOTTOM: *Angle between the spin vectors of the large black holes in each code. The dashed lines indicate the post-merger period.*

different initial data, and yet the behaviour of the apparent horizons that encode the spin information behave identically to a tolerance of ± 0.05 rad between these simulations throughout the entire evolution, if spins are aligned sufficiently well.

Having been suitably aligned the simulations used in this study are summarised in TABLE 5.4, which provides a brief description of their state at t_{ref}^d. Note that the quoted uncertainty estimates relate to the particular waveform under consideration and are not a generic statement on the code that ran the simulation. Higher resolution versions of the BAM and SpEC simulations have also been run and are labelled 'BAMHR' and 'SpECHR'. These simulations provide an estimate of the truncation errors for the two codes, an upper-bound to the agreement we can expect when we start to compare between them with any differences in excess of that revealing more fundamental variations between codes. Ideally the results of the inter-code comparisons would be of the same order as the intra-code comparisons, differing only as a result of resolution.

Label	Separation (M)	Spin/separation angle (rad)	$M r_{ext}$	e (10^{-3})	$M\Omega$
BAMHR	11.796	0.288	50	< 2	0.021970
BAM	11.796	0.225	50	< 2	0.021968
LazEv	11.735	1.024	∞	< 4	0.021964
SpECHR	12.578	0.408	∞	< 2	0.021991
SpEC	12.561	0.521	∞	< 0.3	0.022054

TABLE 5.4   *Summary of simulations at t_{ref}^d. Two simulations are included from BAM and SpEC, one at lower and higher resolutions. Separation is that between punctures for BAM and LazEv, while for SpEC it is calculated between the centres of the excised regions. d_{min} is the spatial resolution of the finest refinement level. The use of "∞" indicates that the results were extrapolated to spatial infinity.*

5.3.   *Comparison of the Waveform*

In this section we begin by directly comparing the phase $\phi_h(t)$ evolution of the $\ell = 2$ modes of h, which together with the amplitude $A(t)$ can be used to construct the strain,

$$h_{\ell m}(t) = A(t) e^{-i\phi(t)}.$$

<div align="right">5·15</div>

This has two advantages. First, the phase involves minimal post-processing as it is directly output by all of the codes, and so only has to be packaged up into the strain time series required by the NRII. It is also independent of the detector used so we can be certain that any differences that arise are purely a feature of the simulation.

FIG 5.4 shows the difference in the phase of the (2,2) mode between codes. Immediately it is clear that phase disagreements grow far quicker during merger, accumulating more in the final $50M$ than over the entire evolution up to that point. While the two BAM simulations at different resolutions remain well-synchronised throughout, as may be expected, interestingly the two SpEC simulations differ a little more, particularly close to merger. This is due to a difference in the way these higher resolution versions were set up. The SpEC simulations differ in both their resolution and also in initial data, but use identical initial parameters to generate the data in each case. Those initial parameters are determined using the higher resolution run, which allows for more precise tracking of the orbital phase and spin orientation, and so while the binary may have been aligned perfectly according to our requirements at the reference frequency in the higher resolution run, with lower resolution it may be offset to some degree causing the difference we're seeing here. The same would be true of the two BAM simulations too, but in anticipation of this discrepancy, the initial parameters were instead tuned to preserve the alignment in both the high and low resolution runs. In short, the two BAM simulations have different initial parameters but record approximately the same physical configuration at the reference frequency, while the two SpEC simulations have identical initial parameters but at the cost of a different configuration. Either way, the phasing between them does not exceed π radians throughout the entire inspiral. More informative is the dephasing between simulations from *different* codes. We see that BAM and LazEv drift by just $\frac{\pi}{4}$ radians by $2100M$ and throughout the difference remains of a similar magnitude to the two SpEC simulations, while BAM and SpEC pass this milestone over $500M$ earlier.

The large variance here between BAM and SpEC is rather surprising, particularly given how well the BAM/LazEv phases agree. It is worth taking some time here to explore this. The first explanation that comes to mind is that the alignment procedure used in this study could be flawed. In the ideal case it relies purely on locating $M\Omega = 0.022$ and so should be relatively robust, but as we've seen eccentricity complicates this and leads to misalignment of the binaries of approximately $\Delta t = 30M$. This issue was most
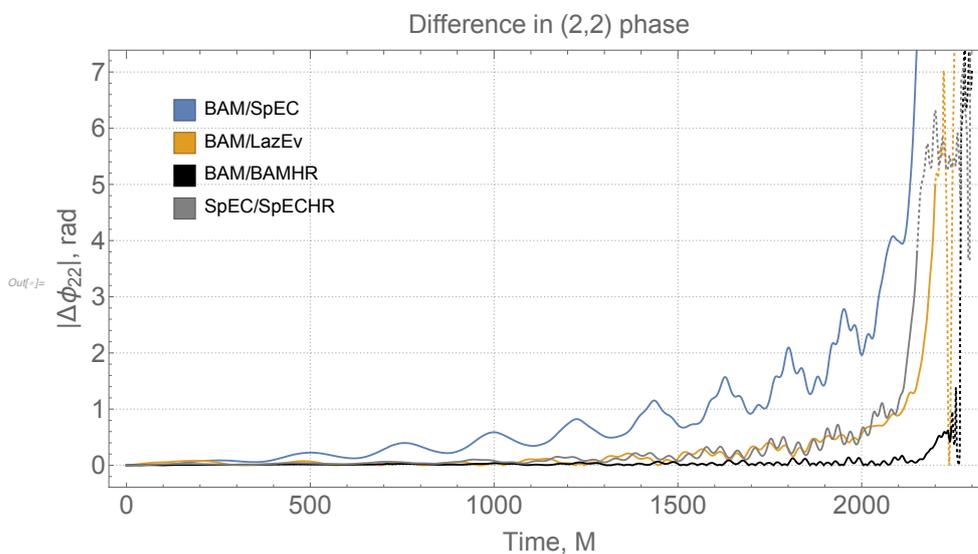
FIGURE 5.4 *Accumulated differences in the phase of the (2,2) mode between BAM and each of the other codes. A shift has been applied such that the phases agree at $t_{ref}^d \equiv t = 0$. The black line shows a comparison of two BAM waveforms of different resolution while the grey shows a similar comparison between high- and low-resolution SpEC waveforms. As these waveforms merge at different times with phase varying rapidly after merger it would be misleading to take their differences when one is merging and the other is still inspiralling. With this in mind we transition from solid to dashed lines to signify that at least one of the simulations has merged.*

pronounced in the LazEv simulation (FIG 5.9), which possesses the largest eccentricity of the five, yet the BAM/LazEv phase difference is dwarfed by BAM/SpEC .

Another possible explanation is that this is a result of the finite radius at which the BAM waveform is extracted, as opposed to the approximation to the waveform as it would be extracted at infinity used by SpEC , described in §5.1. A comparison between them would therefore be expected to reveal this finite-extraction error. This has been quantified for the BAM simulations by constructing the extrapolated waveform. As BAM outputs waveform data extracted at a series of radii, the $\ell = 2$ modes at each of these were aligned about merger (defined as the time when the (2,2) amplitude reaches its maximum), re-sampled, rotated into the quadrupole-aligned frame, and at each time a fit of the form $A + \frac{B}{r}$ was made to the phase according to the various radii. The limit of this fit as $r \rightarrow \infty$ is then taken, with the final extrapolated waveform consisting of the limit of each of these fits at each time step. The difference between the BAM phase and this extrapolated BAM phase is shown in FIG 5.5, where it's clear that this error cannot account
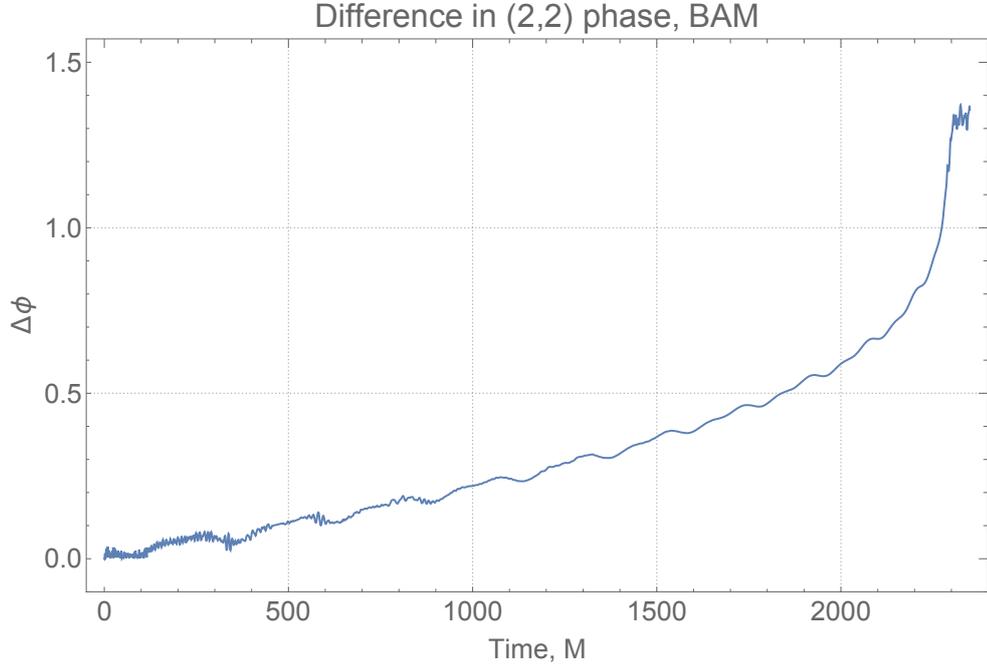
FIGURE 5.5   *Phase difference between the (2,2) modes of the BAM simulation in the QA frame, extracted at a radius of $50M$ and extrapolated out to ∞. The extrapolated waveform has been computed using gravitational wave data extracted at $r_{ext} = \{50, 60, 70, 80, 90, 100, 110, 120\}\ M$.*

for the dephasing between BAM and SpEC. This is not entirely surprising as LazEv also reports the extrapolated waveform—if the use of finite extraction radius were responsible, it would also be seen there. Interestingly the error is of a similar magnitude as the BAM/LazEv dephasing, and so may well be the cause in that case.

Different scalings of the simulation time between codes would also produce a significant dephasing. All of these simulations are scaled by the total binary mass, so if two codes differed in total mass the times and the distances in identical simulations would differ. Of course this would all be entirely consistent *within* each code, and would report a total mass scaled to 1 and component masses 0.25 and 0.75 accordingly, but when it comes to comparing *between* codes such a difference in scaling would emerge. To illustrate this, we take two copies of the same phase data and introduce in one of them an artificial rescaling of the times. We then perform the same kind of alignment that is used in this study and see that these two phases do increasingly drift apart, just as BAM and SpEC appear to (FIG 5.6). Ostensibly both BAM and SpEC compute the total mass of the

binary using the Christodoulou mass [30, 79], with the irreducible mass of each black hole defined by the area of its horizon. But it is worth noting that rather than calculating the Christodoulou mass itself BAM computes the mass as the sum of the ADM masses of the individual black holes and notes that this has been seen to agree within numerical uncertainty with the Christodoulou mass.

Variations in the total mass could also originate from the different gauge choices between the codes given that the masses of the black holes are calculated on the apparent horizons (§2.4), which are highly gauge-dependent. The two BAM simulations are of course identical in their choices of gauge, as are the two SpEC simulations, so if such a scaling discrepancy exists there would be no contribution to the phasing in these cases. And indeed we see quite good agreement in their phases. The LazEv simulation is particularly interesting because now we have a different code that nonetheless has a difference in the relative phasing of the same order as the BAM/BAMHR and SpEC /SpECHR differences. But recall that BAM and LazEv share much of their underlying framework and methodology, and in particular they both use the 1+log gauge, so if the choice of gauge really was contributing it wouldn't be evident in BAM/LazEv any more than it would in BAM/BAMHR. That leaves us with BAM/SpEC, which certainly *do* have different gauges, and indeed we see a much greater phase difference. It should then be possible to apply some artificial rescaling of the times in our SpEC data that acts to improve the phasing with respect BAM, and examine whether it is of a magnitude that could reasonably come from a difference in gauge choices. In FIG 5.7 we reproduce the BAM/SpEC phase difference but artificially scale the SpEC times. We see that a rescaling of 1% can bring the dephasing in line with that in the other comparisons, so clearly this level of difference in mass definitions can reproduce the phase difference we're seeing. Conversely if the gauge differences are shown unequivocally to be the source of the dephasing, then the use of these two waveforms to analyse the same detected gravitational wave would introduce a 1% systematic uncertainty in the inferred masses of the progenitor system.

It might well be asked whether a similar issue plaguing the spins could contribute as well. Given that the spins are also calculated on the horizon this could certainly be the case, but considering the scale we've seen such differences would be, of around 1%, together with the fact that the spins are almost entirely in the plane, any impact on the phase would be negligible. That's not to say the effect of the spin can be entirely neglected: certainly in cases where it is aligned or anti-aligned with the orbital angular momentum

Difference in (2,2) phase with different mass scalings
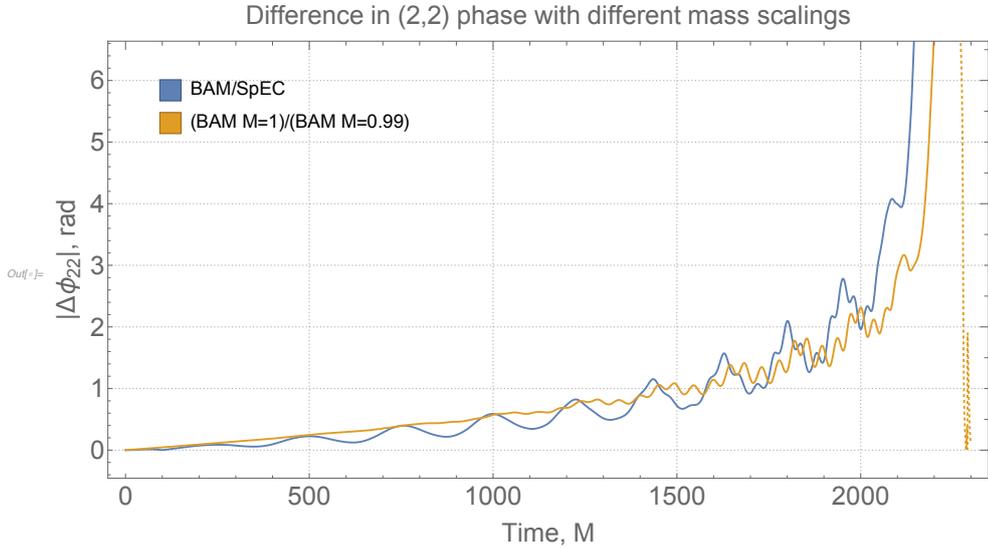
FIGURE 5.6   *In blue the BAM/SpEC (2,2) phase difference. In yellow that of the BAM simulation and the identical data but with the times scaled by a total mass of M = 0.99 rather than M = 1.*

merger occurs sooner and later respectively as a result of spin-orbit coupling [115]. But the spins here are entirely within the plane at t_{ref}^d and throughout the entire evolution don't evolve further than 0.03 rad out of the plane (see FIG 5.3). Clearly then it is the phase evolution that dominates any variation between these codes. It is worth noting that this is not an issue with wave extraction and the gravitational wave output but fundamentally with the evolution of the binary given that a similar degree of variance is seen in both the orbital and (2,2) phases. From the preceding discussion of the (2,2) phase this appears not to be a consequence of precession in and of itself, but rather some as-yet undiagnosed issue with the definition of mass used in the codes producing an inconsistent overall scaling.

Given the significant dephasing that such a discrepancy produces it is perhaps surprising that it hasn't been noted previously. In [90] we see another example of the phase difference between a moving puncture code and SpEC being calculated, found to remain below 0.01 rad at all times. Any problematic definitions or gauge choices, if they exist, would have been present in that study. The approach to aligning the two waveforms is however significantly different between the two studies. In the present study we align them during inspiral and apply a phase shift at that point, while in [90] they are aligned at
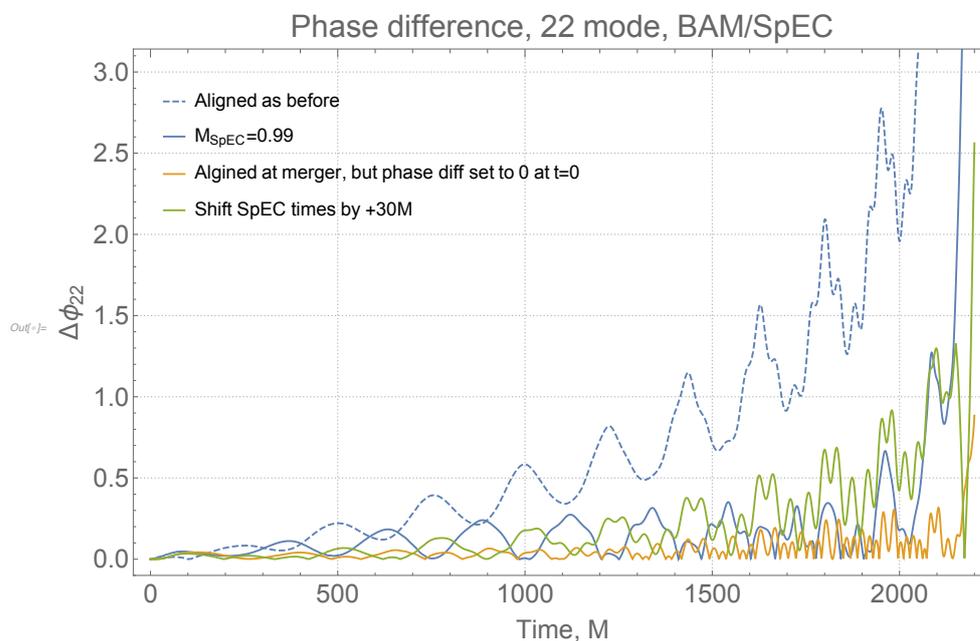
FIGURE 5.7   *The difference in (2,2) phase between BAM and SpEC. The dashed line shows the phase difference in FIG 5.4, the solid blue has had times rescaled to simulate the effect of different overall scalings of the system by total mass. In yellow is the phase difference if, rather than aligning about t_{ref}^d, we align the times at the moment of maximum emission in the (2,2) mode (merger) and apply a phase shift such that $\Delta\phi(0) = 0$, as has been done in [90]. The green line shows the phase difference if the SpEC times are shifted by 30M.*

merger (maximum amplitude in the (2,2) mode), with a phase shift applied at the start of the waveform. Aligning at merger will necessarily lead to smaller accumulated phase differences, and if we follow the same approach here we do indeed see the BAM/SpEC phase difference drastically reduced to < 0.4 rad (FIG 5.7). But as noted in §5.2 an alignment of this kind serves only to suppress differences between waveforms rather than reveal them. To see if precession really is responsible for this phase difference or if it has manifested in prior studies it would be valuable in future work to repeat the SpEC/LazEv comparison of [90] but aligned at some point during early inspiral rather than at merger, or equivalently repeat the analysis in this chapter with an aligned-spin binary.

5.4.   *Comparison of the Precession Dynamics*

So far we've seen how, alongside a contribution from poor extraction of the higher modes, it is the precession dynamics that are primarily responsible for the poor match between our BAM and SpEC simulations. These dynamics are most clearly encoded in the three precession angles, introduced in §5.2. The evolution of the precession angles α, β, γ in each of our simulations are shown in FIG 5.8, from which we can see that, aside from a shift of 2π in the LazEv simulation as a result of choice of convention, there is broad overall agreement. We see in β a variation of no greater than 0.1 rad throughout inspiral, and in α and γ a constant shift of no more than 0.5 rad which can be attributed to the initial misalignment in spins.

Looking to the motion of the orbital planes, in FIG 5.9 we have plotted the direction of the orbital angular momentum vector in each simulation both before and after the spins are aligned at t_{ref}^d. The orbital planes differ in orientation only by 0.02 rad even while undergoing strong precession, so it is clear that the variation in precession dynamics between the simulations is dominated by differences in the initial orientation of the spin.

We see in the orbital phase (FIG 5.10) the same disparity between BAM and SpEC that plagues the gravitational wave phase. One potential source of misalignment in these waveforms could be in the translation from the reference time found in the dynamics data to the corresponding time in the waveform data. We found in §5.2 that a further $\sim 30M$ alignment uncertainty is introduced by this, but this would only affect the waveform data—if we examine the orbital phase we clearly see that the discrepancy persists, even when the possibility of that particular misalignment is eliminated. We can in fact address this concern more generally by considering the integrated phase difference in our waveforms over some range of times, identifying the time shift Δt that minimises the quantity

$$\int_A^B \left| \left(\phi_{orb}^{BAM}(t) - \phi_{orb}^{BAM}(0) \right) - \left(\phi_{orb}^{SpEC}(t + \Delta t) - \phi_{orb}^{SpEC}(0 + \Delta t) \right) \right| dt . \qquad 5.16$$

This time shift is computed over a selection of windows $t = \{A, B\}$, applied to the SpEC waveform in each case, and the resulting phase difference with the BAM waveform shown in FIG 5.11, along with the results for BAM/LazEv and BAM/BAMHR. From this we
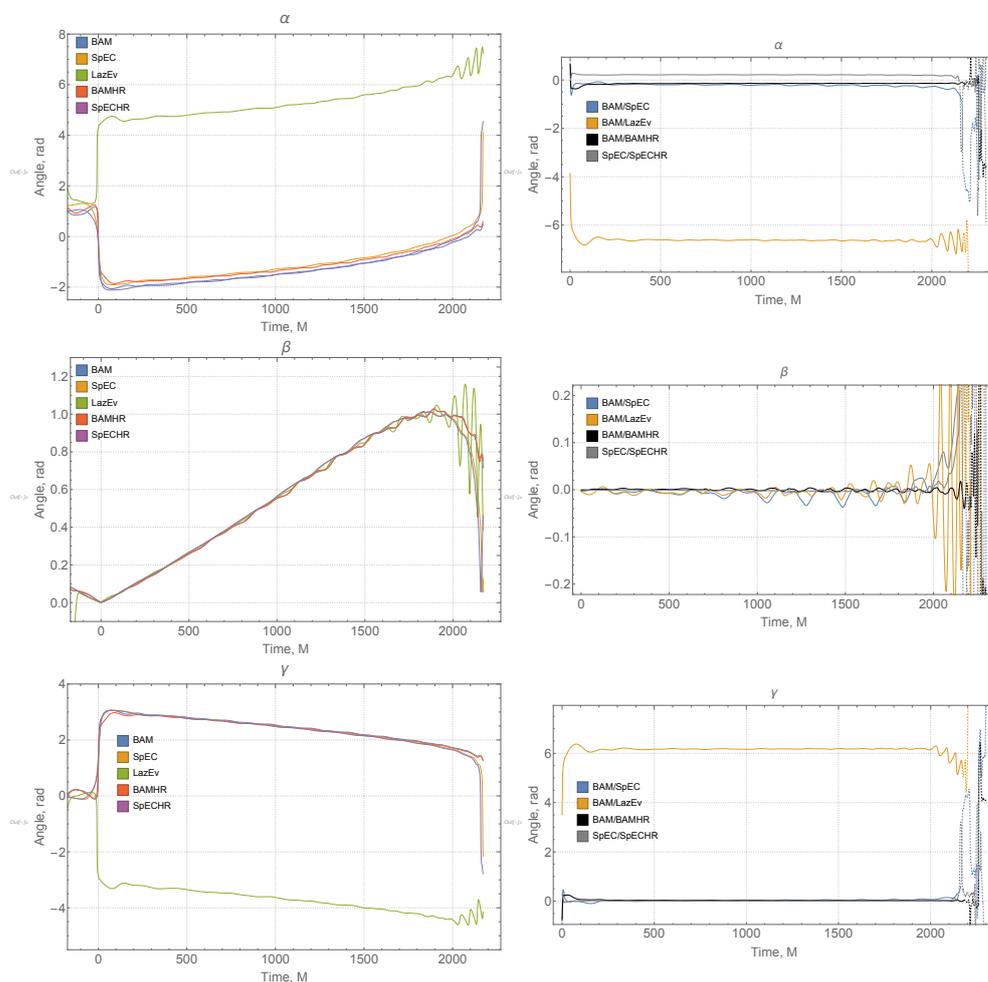
FIGURE 5.8   LEFT: *The precession angles α, β, γ, defined in §5.2.* RIGHT: *the difference in the precession angles between each code.*

see that there is no such time shift that will significantly reduce this difference, until of course the magnitude of the shift is such that we approach aligning at merger, which we are avoiding.

5.5. *Matches*

We also compute the match between our numerical waveforms, defined in §3.6. To give some perspective, the match between a numerical simulation approximating the
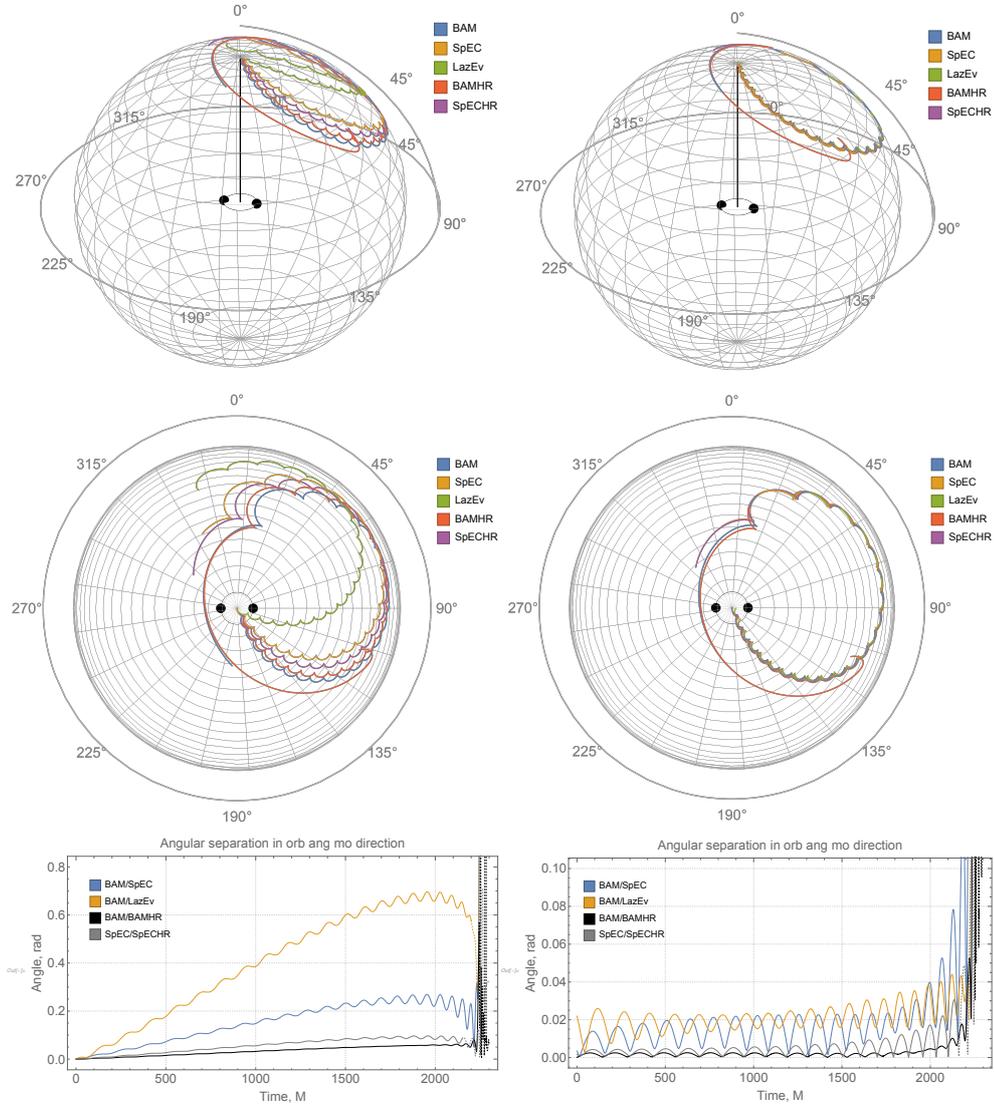
FIGURE 5.9   *The tracks of the orbital angular momentum in each simulation, with phases aligned at t_{ref}^d in the left-hand column and spins aligned at t_{ref}^d in the right-hand column. The two black spheres in the centre joined by a circle represent the black holes and their orbital plane at t_{ref}^d.* TOP: *3D representation of precession tracks on the sky. The solid black line is the vector at t_{ref}^d and is common across all five simulations due to the alignment procedure described in §5.2.* CENTRE: *Top-down view of the same.* BOTTOM: *Angle between the orbital angular momentum vectors in each code. As in* FIG 5.3, *the dashed lines once again define post-merger period.*

Difference in orbital phase
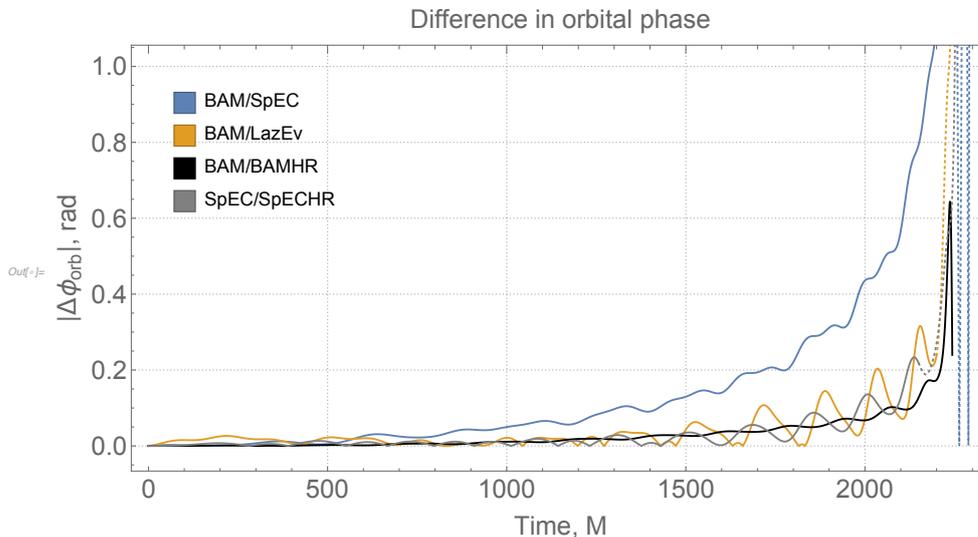
FIGURE 5.10 *Accumulated differences in the orbital phase between BAM and each of the other codes. A shift has been applied such that the phases agree at t = 0. The black line shows a comparison of two BAM waveforms of different resolution while the grey shows a similar comparison between high- and low-resolution SpEC waveforms. As before, dashed lines indicate at least one of the binaries has merged.*

GW150914 event and the waveform model SEOBNRv2 is approximately 0.997 [90]. For our purposes we compute the matches between the BAM waveform and that produced by each of our other codes, and also include on these plots the match between the two BAM simulations of different resolutions and the two SpEC waveforms. The precessing matches here are given as a function of inclination of the orbital plane, optimised over phase and polarisation[3]. These are performed for a starting frequency of 22 Hz with total mass $223M_\odot$ so that the signal starts within the LIGO sensitivity band and are shown in FIG 5.12, with matches performed over all the $\ell = \{2, 3, 4\}$ modes. The detector noise power spectral density used in the match calculations is provided by [116]. We see that, as may be expected, the matches between simulations produced by the same code perform very well, with the SpEC matches in excess of 0.997 at all inclinations and for BAM above 0.999. BAM/LazEv isn't quite as good, but it does for the most part remain above 0.985. The BAM/SpEC match on the other hand displays some concerning features and a particularly strong inclination dependence, at its best rising as high as the BAM/BAMHR match when the binary is face-on (when the (2,2) mode dominates) and at its

---

[3]For an in-depth overview of precessing matches calculations, see Appendix B of [93]
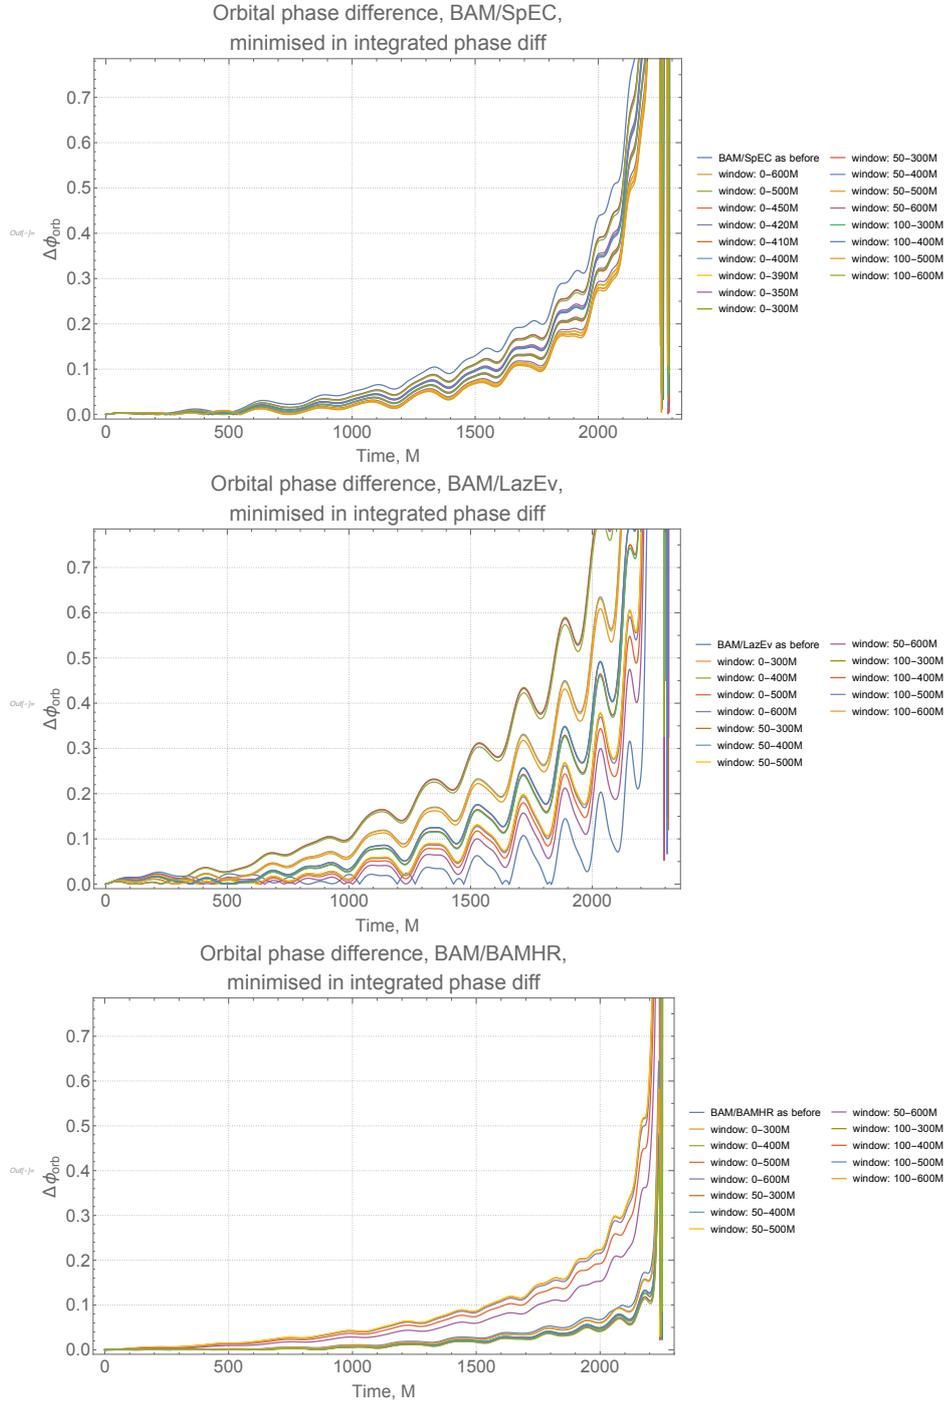
FIGURE 5.11 *Phase differences between BAM/SpEC (top), BAM/LazEv (centre), and BAM/BAMHR (bottom), after applying the time shift that minimises the quantity 5.16, the integrated phase difference, over a selection of windowing regions.*

worst dropping below 0.9. The match averaged over inclination is approximately 0.96. The banks of waveform templates used in LVK search pipelines are constructed such that the minimum match between two neighbouring waveforms in the parameter space is 0.97, corresponding to the 90% confidence region for a signal-to-noise ratio $\rho = 9$ in the LIGO detectors [117].

But what is the reason for such a low match? Given that previous studies [89] have shown matches in excess of 0.999 between non-precessing waveforms of different NR codes, there are two possibilities. Either the code isn't handling the dynamics of the precession accurately, or the effects of precession on the waveform aren't being picked up by the code sufficiently well. Now, in truth there will be some element of both occurring but we can determine which is dominant by splitting our analysis up into two parts. We've seen that the precessing match is poor between our waveforms, so to diagnose the problem we split it into a non-precessing match plus a set of precession angles. It is worth at this point taking a moment to examine how this helps us.

If the precession dynamics are at fault then the code will produce a perfectly valid waveform but one that doesn't correspond to the scenario we thought we were simulating. Specifically, there will be more power in some modes than we would expect and correspondingly less in others and so the match with a 'correct' waveform will be poor—because they are from two different binaries. If this waveform is then rotated into the quadrupole-aligned frame with the $(2, \pm 2)$ modes dominating the spectrum and the waveform made effectively non-precessing, the match should improve dramatically. By rotating into the QA frame we remove the precession from the waveform and so remove any consequence that could arise from inaccurate precession dynamics. The match will now be between two waveforms generated by identical binaries, and the cause of the poor precessing match lies in the precession dynamics.

In the second case, the binary precesses perfectly but the changes wrought on the waveform as a result of this are not picked up sufficiently well. Poor extraction of the higher modes will produce a waveform that doesn't exactly correspond to the dynamics of the binary. As we've seen in §5.2, a precessing binary displays mode-mixing, a shifting of power from the dominant $(2, \pm 2)$ mode into the other, usually subdominant, modes. We can imagine that, say, the $(3, 2)$ mode isn't extracted accurately by one code and so some power is lost from it, and would as a result produce a poor match against a 'correct'
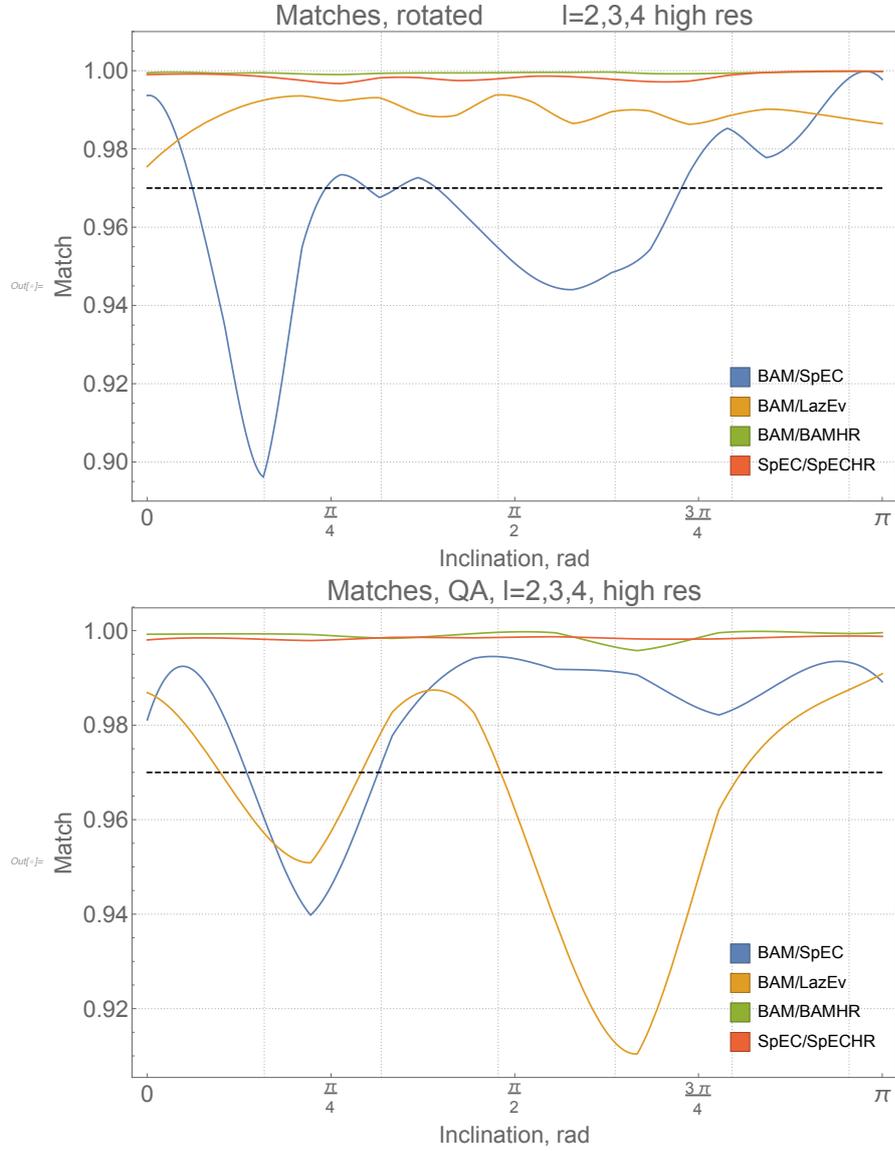
FIGURE 5.12   *The match between each of the waveforms, where in the right-hand panel they have been rotated into the quadrupole-aligned frame. For a signal of $\rho = 9$ the 90% confidence region in LIGO corresponds to a match of 0.97, shown by the dashed black line.*

waveform. This time however if we rotate into the QA frame the match will remain poor, as the lost power won't be restored and the waveform will remain inaccurate.

The results of the matches between the quadrupole-aligned waveforms are shown in the bottom panel of FIG 5.12, and between BAM and SpEC we do indeed see an improvement, suggesting that there is a greater impact from precession dynamics than from the extraction of higher modes. That's not to say the modes have no impact: if we plot the match between individual modes of the BAM and SpEC waveform (FIG 5.13) we see a much more nuanced story starting to emerge. The matches between the $m = 0$ modes are consistently poor, only surpassed by the $(4, \pm4)$ modes. For a non-precessing binary it is expected that the match between these higher modes would be relatively poor given how little power they contain with numerical noise washing out much of the detail, but in a precessing case we see each mode modulated in power throughout the evolution, each with times when it is particularly strong and times when it is weak enough to be significantly affected by noise. This would therefore be expected to impact each of the modes equally. The issue then is one of frequency: as we move to higher m the frequency of the signal increases, and so will require finer resolution in time to resolve each cycle to the same accuracy as the (2,2) mode. These higher-m modes then will by definition always be more coarsely resolved than the (2,2) mode and exhibit poorer matches. This is offset somewhat in the $\ell = 2$ matches by the reduced power content in the $m = \pm1$ modes (see FIG 5.14), where noise contributes a greater proportion of the signal. The poor handling of these modes certainly drags the overall match down then to some extent, even if precession dynamics is the dominant cause. Interestingly though, between BAM and LazEv we see the match get *worse* when we move to the quadrupole-aligned frame, suggesting that here it is the higher modes that are the dominant source of discrepancy between these particular simulations. Again turning to the mode-by-mode matches as before (FIG 5.13) the $m = 0$ modes are seen to perform even worse than their counterparts in the BAM/SpEC matches, and the $m = \pm1$ matches are also now significantly lower. FIG 5.14 shows the relative power contained in each of these modes, taking BAMHR in the inertial frame as representative.

It is perhaps more intuitive to convert the matches into the SNR at which our waveforms would be distinguishable from one another in the current generation of gravitational wave detectors. Suppose that two NR waveforms are being used to infer the source parameters of a gravitational wave that has just been detected. If they are pre-
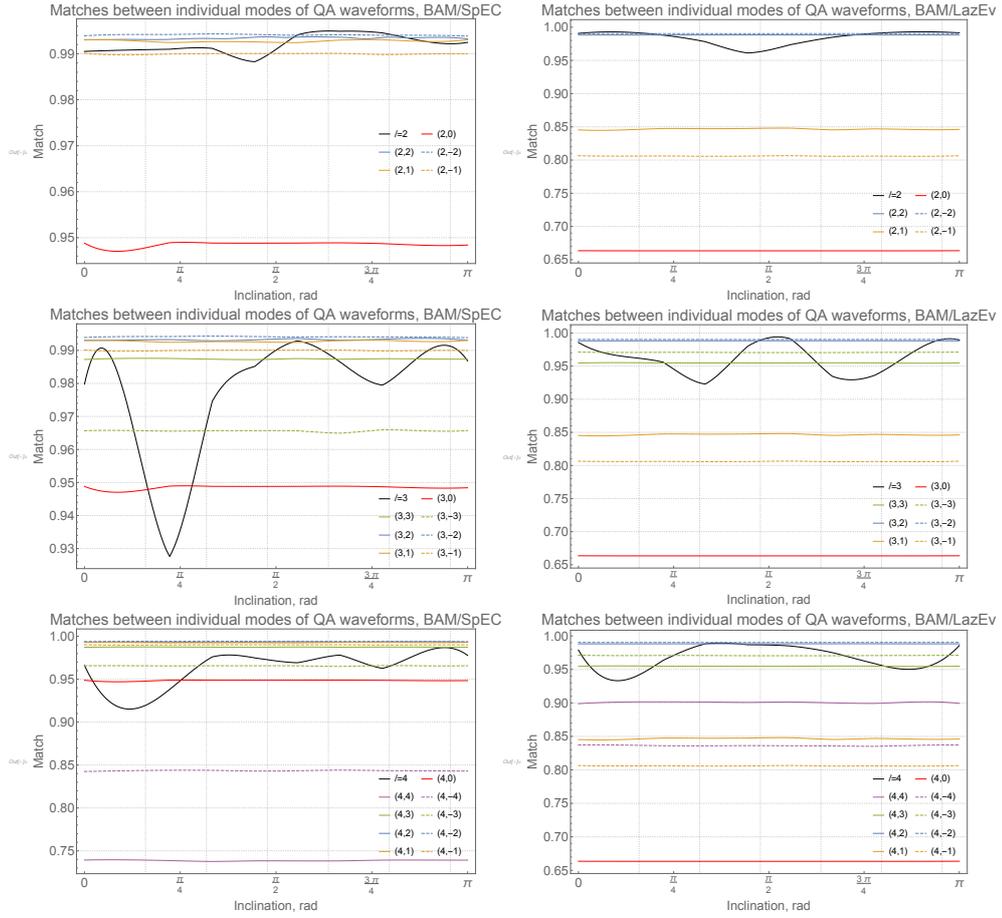
FIGURE 5.13 *Matches computed for each individual mode for BAM/SpEC (above) and BAM/LazEv (below). The top row features the $\ell = 2$ modes, the middle row the $\ell = 3$ modes, and the bottom row $\ell = 4$. In each of these the black line indicates the match computed over all modes in the particular ℓ.*
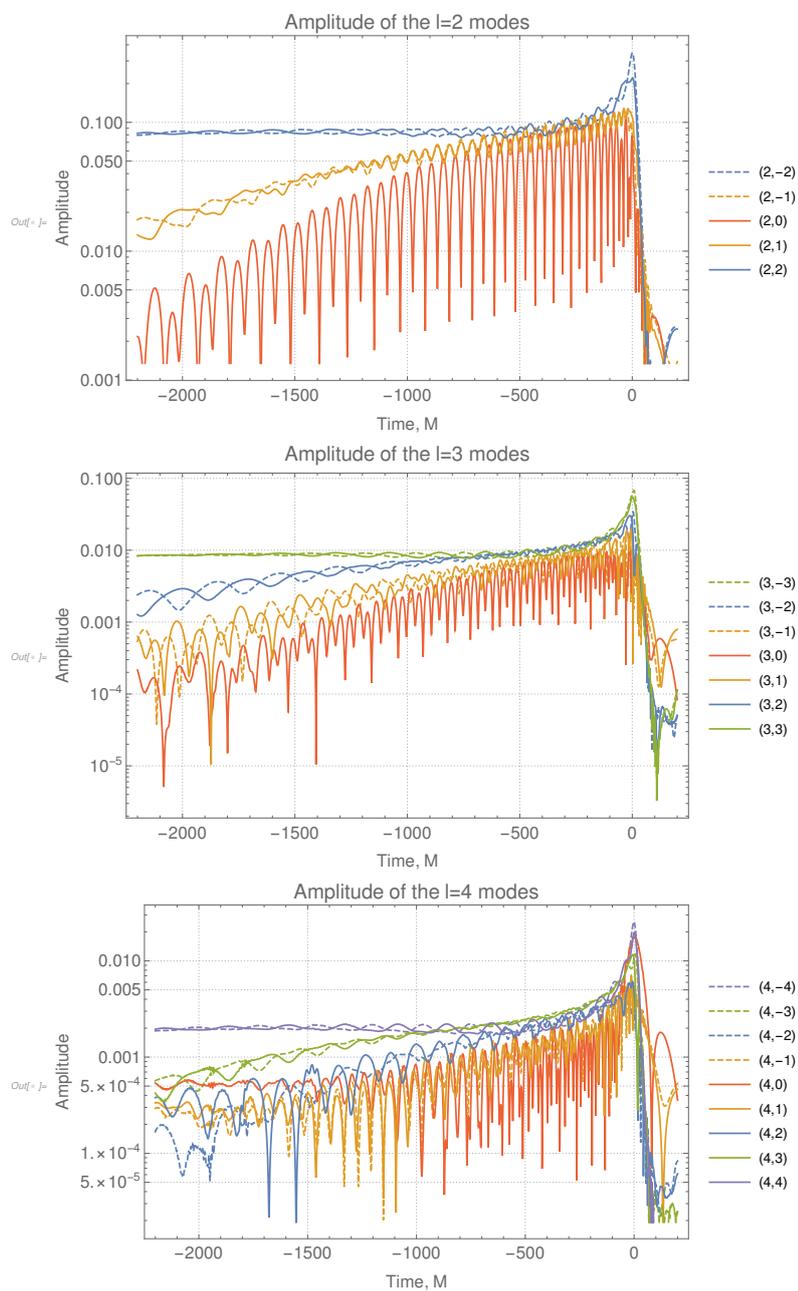
FIGURE 5.14 *Amplitudes of the $\ell = \{2, 3, 4\}$ modes of the BAMHR waveform in the inertial frame. Colour scheme is that used in* **FIG 5.13,** *with dashed lines indicating the negative m modes.*

cisely identical they will return identical estimates of these parameters. If however they differ, naturally they will report different results. But suppose the detected signal were very weak. In that case the differences between the waveforms may be too small to be seen: noise will drown out any fine differences between the two waveforms unless the signal is sufficiently loud. In practice our numerical waveforms will never be perfectly identical simply due to numerical noise so there will always be some finite SNR at which they would be distinguishable. The question, though, is whether there are sources other than numerical noise that produce differences, and whether these differences are of sufficient magnitude to be noticeable in the data from our gravitational wave detectors. As points of reference the typical detection threshold within the LIGO detector is an SNR of 8 [118], the loudest black hole binary detected at time of writing had a network SNR of 24 [3], and while in theory there is no particular upper limit, results from population and event rate studies suggest that, for a random sample of 100 gravitational wave events, there is a 50% probability that the loudest event is greater than 60 [119].

Baird et al. in [117] provide a straightforward conversion between match $\mathcal{M}$ and this critical SNR ρ_c above which two waveforms are distinguishable,

$$\rho_c = \sqrt{\frac{\chi_k^2(1-p)}{2(1-\mathcal{M})}}\,, \tag{5.17}$$

where $\chi_k^2(1-p)$ is the chi-squared value with probability $1-p$ of obtaining that or greater value.

This was calculated in [89] for the (2,2) mode of an equal-mass non-spinning binary, finding $\rho_c \sim 25$ between BAM and SpEC . In FIG 5.15 we plot the critical SNR for the waveforms considered in this study, calculated across all the $\ell = \{2, 3, 4\}$ modes at 90% confidence for $k = 7$ for the seven intrinsic parameters of the binary. For waveforms in the inertial frame we see that while differences between waveforms generated by the same code remain safely beyond the current loudest black hole binary signal at $\rho_c > 30$, the comparisons between different codes is consistently below $\rho_c = 25$. BAM/SpEC performs particularly poorly here, with all except face-on inclinations distinguishable at $\rho \sim 10$.
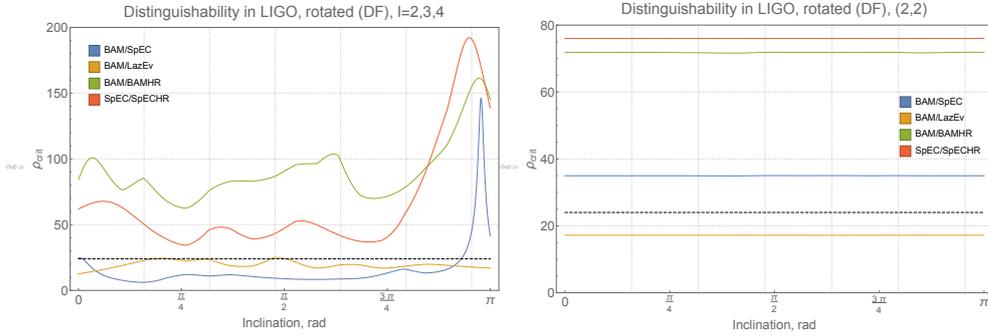
FIGURE 5.15   LEFT: *The distinguishability between each of the waveforms in LIGO (design sensitivity) for the match computed over all $\ell = \{2, 3, 4\}$ modes.* RIGHT: *Distinguishability for the match in the (2,2) mode only. In both panels the dashed black line shows the highest SNR binary black hole detection to date, GW150914.*

5.6.   *Discussion*

Of primary importance moving forward is the investigation of this observed phase difference between the moving puncture and excision codes. A clear and relatively straightforward test would be repeating the SpEC/LazEv comparison in [90] without the alignment at merger, with the waveforms aligned instead around $2000M$ before merger as has been done in this study. This would allow us to define the configuration without the complication of adjusting the spin configuration, and focus on achieving exactly the same horizon quantities and eccentricity in the simulations. It would then be immediately clear whether the phase differences seen here are merely the result of precession dynamics or represent something more fundamental, such as the mass discrepancy discussed above. If the same degree of difference in phases is seen the source will need to be identified. One possibility is the assumption in BAM that the ADM mass of a Kerr puncture sufficiently approximates its Christodoulou mass, which it is claimed to achieve to within numerical accuracy for non-precessing binaries but should be verified for the precessing case. Another is that the issue lies with the gauge choices: damped harmonic for SpEC, 1+log $\widetilde{\Gamma}$-driver for BAM and LazEv. These do after all differ significantly—for example, the 1+log slicing enjoys strong singularity avoidance as slices are unable to approach the singularity, but with harmonic slicing they have been shown to approach it (albeit in infinite coordinate time) [120].

There also remains the discrepancy between higher modes in BAM and LazEv. That these discrepancies persist even when rotated into the QA frame is suggestive of this being an issue in the extraction of Ψ_4, and as such an in-depth study of the differences between their extraction methods would be recommended.

The eccentricity-induced ambiguity in defining a reference frequency produces a non-trivial misalignment of spins, leading to different precession dynamics. This however is entirely an issue with the setting up of simulations for this study rather than something inherent to the codes themselves: given initial data constructed with vanishing eccentricity the orbital planes would precess identically to within a tolerance of 0.03 rad. A case could certainly be made therefore that, given how significant an impact such small uncertainties in the frequency can have, there is need for stricter requirements on NR waveforms than have currently been implemented. The first priority would be some amendment to the NRII is recommended regarding the use of orbit-averaged versus instantaneous orbital frequencies. Currently the orbital frequency dataset to be included in the final data product is that defined by 5.2, but as we have seen this is subject to oscillation due to the eccentricity. Rather than replace this with the orbit-averaged frequency, we would recommend including both the orbit-averaged frequency and instantaneous as separate datasets as both could well prove to be useful for eccentricity studies or the development of eccentric waveform models. Similarly, tighter restrictions on acceptable eccentricities would appear to be necessary given how even small eccentricities have been shown to affect the spin configuration significantly.

Given the magnitude of the impact on phase that has been seen in this study, greater transparency regarding the various conventions used in each simulation is strongly recommended. In the context of detections it is the relative phasing of the waveforms that impacts the match between them, so a potential incompatibility in, say, choices of gauge or definitions of eccentricity between codes risks significantly impacting the science output of the collaboration.

As far as the future is concerned, as we approach the era of third generation detectors such as Einstein Telescope and Cosmic Explorer with their expected SNRs exceeding 100 [8] it is clear that, unless there are drastic improvements in the consistency of NR simulations, waveforms produced by different codes will be readily distinguishable. Results from parameter estimation will differ depending on which code's waveforms are used. Currently numerical waveforms are not directly used in parameter estimation, with a few

numerical waveforms instead produced at various points in the parameter space which are then used to inform the development of models that approximate the key features of the waveforms. Ideally the entire parameter space would be populated by thousands of numerical waveforms to be used directly in parameter estimation as they offer greatest accuracy. As this is prohibitively expensive in terms of both time and computational resources, these models are used instead. It has been shown here that even in the ideal case, of purely numerical parameter estimation with no errors due to the approximations made by the models, results would still differ depending on the choice to use either BAM, SpEC, or LazEv.

The impact of multi-detector networks has been neglected in this work. This is an area ripe for future study in the context of numerical waveform systematics given it can contribute an effective amplification of a signal, as it may well increase the likelihood that an event is detected of sufficient power to fall within the distinguishable region. Additionally as current detectors are improved over their operational lifetime and their range is steadily extended the volume of space open to us grows, with a corresponding increase in the number of potential detectable events. Purely on probabilistic grounds the prospect of an unusually loud signal becomes increasingly likely, together with the prospect of our numerical waveforms returning conflicting parameter estimates as a result.

Part Three

TOWARDS EXTREME
MASS RATIOS

Stable Evolutions in Coarse Resolutions

6

As gravitational wave astronomy advances towards ever-higher sensitivities and as-yet unexplored frequencies, the prospects for the detection of a signal originating from a binary of high mass ratio become ever more hopeful. The waveform models and simulations we have explored in PART I and PART II provide sufficient detection and analysis tools for the current era of second-generation ground-based detectors—mass ratios up to $q = 8$ are receiving a lot of attention within the community, and it is likely that waveform uncertainties will remain beyond the accuracy requirements for observations with second-generation detectors. A number of potentially precessing signals have been detected and analysed using these precessing models, but even then parameter estimation results remain minimally affected by the use of either precessing or non-precessing models. All published black hole binary signals at the time of writing have been approximately equal-mass, save for GW190412 with posterior support for $q = 3$ [87][1]. With a low-frequency limit in the LIGO detectors of 10 Hz (corresponding to the innermost stable circular orbital frequency of a $M = 400 M_\odot$ binary) and an approximate lower mass limit on black holes of $3M$ [121], the highest currently detectable mass ratio is $q \approx 145$. But the upcoming space-based gravitational wave detectors—Taiji, LISA, and TianQin—unimpeded by seismic noise and noise from the mirror suspension systems will be sensitive to far lower frequency gravitational wave signals. The LISA frequency range for example is 0.0001–1 Hz, which will allow for detections of binaries of chirp masses in the range 10^3–$10^7 M_\odot$. As well as mergers of intermediate-mass or supermas-

---

[1] GW190814 has an estimated mass ratio of $q = 9$, but is most likely the result of a neutron star-black hole merger.

sive black holes, this will permit the detection of *extreme mass ratio inspirals* (EMRIs) in which a stellar-mass black hole merges with one of its supermassive counterparts. With gravitational wave signals that lie within the LISA frequency band for the final few years of inspiral [122] we have a way to directly probe the strong-field geometry close to a black hole horizon, as well as access to the rich harmonic structure during ringdown that in the comparable-mass regime is dominated by the $(2, 2)$ mode, together providing a means to test general relativity *in extremis*. As we've seen throughout this thesis, accurate waveform models are essential to gravitational wave astronomy. But to construct these we need accurate predictions from theory as to the phenomenology of these waveforms; predictions that can only be generated by numerical relativity. And while current numerical relativity codes are perfectly capable of simulating such high mass ratios, the timescales even a single simulation would require to run for sufficient duration are entirely impractical. A dense sampling of the mass ratio and spin parameter space is therefore out of reach, and so too then is an accurate waveform model.

To see why this should be the case we first note that the black hole solution is asymptotically Minkowskian, with the BSSN variables in our simulations taking on constant values in the far-field limit. It is only close to the puncture that the spacetime—and correspondingly the BSSN variables—changes significantly over short length scales. Clearly then in order to accurately capture the solution on the numerical grid a finer resolution is required in the region close to the puncture than would be needed further away, and it is precisely this that motivated the use of mesh refinement in numerical relativity codes 30 years ago [123], that is, the use of a set of nested numerical grids of increasing resolution and decreasing extent, with the finest grids centred on the puncture (see §2.5). This has been used successfully in the production of black hole binary simulations ever since and as we've seen has been invaluable in producing waveforms used in the search for and analysis of astrophysical gravitational wave events in the LIGO, Virgo, and KAGRA detectors. But even now these simulations are prohibitively expensive to run for more than a dozen orbits before merger, to cover the parameter space systematically and satisfactorily, or to simulate binaries of mass ratios greater than approximately $q = 20$. The reason for this is the scaling of these simulations by the total mass of the system, M. Increasing the mass ratio is therefore equivalent to reducing the size of one of the component black holes, and as a result increasingly fine resolutions are required to accurately capture the smaller component, without which the solution relaxes to flat space (or, more colloquially,

"*evaporates*"). Not only must the spatial resolution increase, but the temporal resolution too: typically a Runge-Kutta evolution with finest grid spacing d requires time steps of a size no greater than $\frac{d}{2}$ to evolve stably [127], so as finer and finer grid spacings are required at higher mass ratios the number of time steps per orbit must correspondingly increase linearly. The memory required by the simulation is also subject to scaling as resolution is increased, and so for full 3D simulations the computational cost scales as the cube of the mass ratio.

It is instructive to consider an example that demonstrates this explicitly. A general rule-of-thumb is that 50 grid points covering the black hole is sufficient for accurate numerical evolutions (see for example [128–130]). In isotropic coordinates the horizon radius is $\approx 0.7m$, leaving us with a grid spacing of $d = 0.028m = 0.014M$ in the case of an equal-mass binary which as we have just seen requires time steps of at least $\Delta t = 0.007M$. The orbital period of the binary explored in CHAPTER 5 was approximately $300M$. Taking this as a representative example we see that the ten orbits of inspiral typical of the simulations used in gravitational astronomy equate to 430,000 time steps. The equal-mass simulations produced in CHAPTER 4 completed ten orbits in one week on 256 cores. If we now consider a $q = 100$ binary the above calculation leads to 2.16×10^7 time steps—a year of continuous evolution. And this would represent only the extreme lower end of the EMRI regime; with LISA and other space-based gravitational wave detectors we expect to see mergers of stellar mass black holes with supermassive black holes—mass ratios of up to 10^9—with signals that would be in-band for thousands of orbits before merger. Using the same kind of back-of-the-envelope calculations, evolving a $q = 10^6$ binary for 100 orbits would take 100,000 years.

There have nonetheless been successful efforts to simulate mass ratios higher than is typical: BAM for example has been used to evolve a $q = 18$ binary [70], and RIT have performed $q = 100$ binary evolutions using LazEv [124, 125]. But the BAM effort, despite possessing a mass ratio that is near negligible on the EMRI scale, still required over 100,000 CPU hours per orbit. The earlier LazEv simulation ran for just two orbits—insufficient for use in searches and parameter estimation, which require dozens of gravitational wave cycles [126]—was non-spinning, and still required a month of wall-clock time per orbit at a cost of 250,000 CPU hours, even with a temporal resolution chosen to minimise expense at the cost of significant mass loss in the black holes. The more recent simulation required 18,480,000 CPU hours to complete 13 orbits. It may therefore be

regarded as a valuable demonstration that there is no fundamental obstacle to these high mass ratio simulations beside computational cost.

The issue then is entirely one of resolution, not some more fundamental inability or lack of understanding. If a way were found to stably and accurately evolve a black hole in extremely coarse resolutions the whole issue could be sidestepped. This chapter aims to take the first steps towards this with the introduction of a novel approach to the numerical evolution of black hole spacetimes, and begin a programme of work that it is hoped will result, in time, with long-term evolutions of binaries of arbitrarily high mass ratio. We have conducted, and here outline, proof-of-principle tests on a single Schwarzschild black hole, a boosted non-spinning black hole, and finally the head-on collision of two non-spinning black holes of mass ratio 100, and have demonstrated that the method is stable. Stable, accurate evolutions of the Schwarzschild spacetime were achieved for a duration of $30{,}000M$; a boosted Schwarzschild black hole with boost velocity 1% the speed of light has been run until destabilisation at $3500M$; and promising initial results with the $q = 100$ head-on collision of two non-spinning black holes have been achieved, all three cases with resolutions approximately three orders of magnitude more coarse than previously possible.

6.1. *A New Approach*

Currently numerical evolutions of black hole binaries solve the Einstein equations at every grid point at every time step. In this work we aim to reduce the computational load by taking advantage of the fact that a great deal of analytic information is known regarding the form of the solution near the black hole. For a given (non-spinning) binary, we know that locally the spacetime around each black hole is described by the Schwarzschild solution—of course as you move further out this is modified due to the presence of the other black hole, but for a sufficiently high mass ratio it is hoped that these modifications are characterised by a length scale large enough that they can be treated as a smooth, slowly-varying background to the foreground small black hole, determined by the length scale of the larger black hole. We can therefore, for some solution variable $f(x^i, t)$, decompose into the local solution at the small black hole $f_0(x)$ and the background due
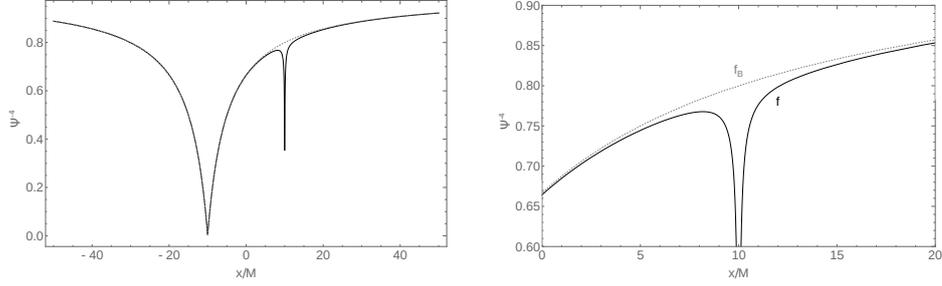
FIGURE 6.1   *The conformal factor to the negative-4th power for data representing two initially station-ary black holes of mass ratio q = 100, located at x = ±10M. The solid line shows the full solution $f(x, t)$, and the dashed line is our background $f_B(x, t)$. The left panel shows both functions for both black holes, while the right-hand panel zooms in on the region close to the small black hole.*

to the larger black hole $f_B(x, t)$:

$$f(x, t) = f_B(x, t) + f_0(x) . \qquad\qquad 6·1$$

FIG 6.1 illustrates this for the inverse fourth power of the Brill-Lindquist conformal factor in a q = 100 binary, with the black holes separated by 20M. The solid line shows the full solution $f(x, t) = \psi_{BL}^{-4}$, and the dashed line is our background $f_B(x, t)$. Note that this is an illustrative example for the purposes of introducing the concept only—Brill-Lindquist wormhole data will not be used, and if it were the form shown in FIG 6.1 would not be retained beyond the initial time step. From the introduction of Brill-Lindquist data in PART I we have the conformal factor for a multiple black hole spacetime given by 2·5

$$\psi^{-4} = \left(1 + \frac{m_1}{2r_1} + \frac{m_2}{2r_2}\right)^{-4} , \qquad\qquad 6·2$$

where

$$r_i \equiv \sqrt{(x - x_i)^2 + y^2 + z^2} . \qquad\qquad 6·3$$

This corresponds to the black line in FIG 6.1. The grey line is simply the conformal factor of the large black hole alone,

$$\psi_2^{-4} = \left(1 + \frac{m_2}{2r_2}\right)^{-4} , \qquad\qquad 6·4$$

which lacks the divergent term due to the smaller black hole. Clearly a far coarser grid spacing will suffice to resolve $f_B(x, t)$ and its derivatives than that required by $f(x, t)$, and as $f_0(x)$ is known *analytically* its derivatives can be calculated to arbitrary precision. By treating these contributions separately the derivatives of the global solution variable $f(x, t)$ can be computed with a far larger grid spacing, or in other words a far coarser resolution, than is currently required. It should be possible then to take advantage of this by replacing each of our evolution variables with their decomposed forms—the analytic local solution plus some correction term—and numerically evolve only this correction. Accordingly throughout these chapters we label these the *analytic piece* and the *numerical piece*, which together make up the full variable. Our evolution variable f is now decomposed as:

$$
\begin{aligned}
f &= f_{analytic} + f_{numerical} \text{ ,} \\
&\equiv \mathbb{f} + \boldsymbol{f} \text{ .}
\end{aligned}
$$

<div align="right">6·5</div>

where for the example shown in FIG 6.1 the analytic piece is the local solution at the small black hole $f_0(x, t)$ while the evolved numerical piece is the background $f_B(x, t)$. The somewhat idiosyncratic notation here has been chosen so as to express the distinction between the three quantities (full, analytic, numerical) while preserving the notion that they all refer to the same solution variable without further proliferation of sub- and superscripts that risk being mistaken for tensorial indices. It is hoped that this choice assists in a more intuitive reading, suggestive of the role of the analytic quantities as a sketched outline to be filled in by the numerical piece.

Throughout this work our evolution variables are the lapse α, shift vector β^i, conformal factor $\phi = \log \psi$, and the extrinsic curvature split into its trace K and trace-free part A_{ij}. These are then decomposed as

$$
\begin{aligned}
\alpha &= \mathbb{a} + \boldsymbol{a} \text{ ,} \\
K &= \mathbb{K} + \boldsymbol{K} \text{ ,} \\
\phi &= \mathbb{\phi} + \boldsymbol{\varphi} \text{ ,} \\
\beta^i &= \mathbb{\beta}^i + \boldsymbol{\beta}^i \text{ ,} \\
A_{ij} &= \mathbb{A}_{ij} + \boldsymbol{A}_{ij} \text{ .}
\end{aligned}
$$

<div align="right">6·6</div>

In practice this approach involves a simple reformulation of the evolution equations in which the derivatives become a mix of these analytic and numerical pieces and their derivatives. The arbitrary precision to which these analytic derivatives can be calculated is the key feature of the new approach we outline here that allows for coarse-grid evolutions.

At the start of every time step we redefine each evolution variable $f \rightarrow f - \mathring{f} = \bar{f}$. This redefinition means that the evolution subroutine must be adapted accordingly: everywhere f appears will now be referring to $\bar{f}$, so must be replaced with $\bar{f} + \mathring{f}$ to rebuild to the full quantity for evolution. This will require the first and second derivatives of $\mathring{f}$. Finally after the evolution step is complete we will need to return to the full variable, so we again redefine every instance of f (which we recall at this stage is currently referring to $\bar{f}$) as $\bar{f} + \mathring{f}$. And so we end the time step with the full f intact. The only other change that needs to be made is in the boundary conditions on the lapse, as it is the only one of our quantities that does not vanish in the far field limit. As a result, after subtracting the analytic piece the leftover numerical piece *will* vanish at the boundary, so the conditions need to be changed accordingly.

In these preliminary tests we supply the analytic piece in the form of simple polynomial fits to highly accurate spherically symmetric numerical data representing the Schwarzschild solution in trumpet form (see §2.3) for each evolution variable. Producing such a fit has been achieved through the use of Padé approximants applied to high-resolution static trumpet initial data, computed in [34] as part of efforts to incorporate the trumpet topology into BAM simulations, and informed by knowledge of both the near-puncture and asymptotic behaviour. This data is provided for α, K, ψ, and β_r, for the radial coordinate r. Given the spherical symmetry of the system the radial shift vector data is used to construct the fit $\mathring{\beta}^r$, which is then transformed into the Cartesian form required by BAM by simply

$$\mathring{\beta}_i = \frac{\partial r}{\partial x_i} \mathring{\beta}_r \ . \tag{6.7}$$

The only minor obstacle here comes in the extrinsic curvature. This is not included in the high-resolution trumpet data, but can be constructed from data that is available to us as [10]

$$A_{kl} = \psi^{-4} K_{kl} - \frac{1}{3} g_{kl} K \ . \tag{6.8}$$

The only missing ingredient then is the g_{kl}, but working as we are with conformally flat Schwarzschild coordinates these are trivially the flat space metric components. It is then

straightforward to construct the A_{kl} in spherical coordinates. Then, analogously to the treatment of the shift, a fit is made to these known spherical A_{kl} to which the appropriate tensorial transformation is applied to get it in Cartesian form,

$$\mathbb{A}_{ij} = \frac{\partial x_k}{\partial x_i} \frac{\partial x_l}{\partial x_j} \mathbb{A}_{kl} \, .$$

6·9

The fits that have been constructed in this manner are

$$\alpha = \frac{0.0100r^{0.545} + 0.9311r^6 + 2.9584r^5 + 3.9283r^4 + 1.1093r^3 - 0.0207r^2}{0.0146r^{-0.545} + 0.9311r^6 + 3.8896r^5 + 7.3514r^4 + 6.7549r^3 + 1.9378r^2} \, ,$$

$$\mathbb{K} = \frac{1}{2.3046r^{3.663} + 6.8388r^{2.276} + 1.5983r^{1.218} + 5.7667r^{1.074} + 3.3229} \, ,$$

$$\varphi = \frac{\frac{4.7557}{r^{2.016}} + \frac{4.5876}{r^{2.016}} + \frac{4.5753}{r^{2.016}} + \frac{4.5726}{r^{2.016}} + \frac{4.5442}{r^{2.016}} + \frac{18.3270}{r^{0.964}}}{\frac{8.3174}{r^{1.942}} + \frac{18.6008}{r^{1.637}} + \frac{13.0821}{r^{1.018}} + \frac{15.8568}{r^{0.8889}} + \frac{15.6526}{r^{0.889}} + 36.1443r^{0.038} + \frac{1}{10}} \, ,$$

$$\beta_r = \frac{1}{2.465 + 0.476r^2 + 1.755r + \frac{1.784}{r}} \, ,$$

$$\beta_\theta = 0 \, ,$$

$$\beta_\phi = 0 \, ,$$

$$\mathbb{A}_{rr} = -\frac{0.9}{0.0560r^{3.329} + 0.6593r^{2.416} + 1.5582r^{1.133} + 1} \, ,$$

$$\mathbb{A}_{\theta\theta} = \frac{1}{0.4788r^{1.001} + 3.0123r^{-0.937} + 1.5793r^{0.050} + 0.0001r^2 + 2.2055r^{-2}} \, ,$$

$$\mathbb{A}_{\phi\phi} = \mathbb{A}_{\theta\theta} \sin(\theta) \, ,$$

6·10

and in FIG 6.2 we demonstrate how they compare to the high-resolution trumpet data. While the numerical piece does still possess some sharp features close to the puncture, these are of a magnitude small enough that the effect of their exclusion in the global solution is deemed to be negligible.

One further modification can be made at this point for the sake of future-proofing. Given that the ultimate goal of this effort is to apply these fits to the evolution of the smaller component in a high mass ratio black hole binary, some consideration needs to be given to the scaling of these fits with mass. For the most part this is straightforward: the mass contributes an overall rescaling of the coordinate distance, so we simply modify 6·10 by making the change $\{x, y, z\} \rightarrow \{\frac{x}{m}, \frac{y}{m}, \frac{z}{m}\}$, taking care not to neglect the extra factor of m this will contribute in the derivatives. A further factor of $\frac{1}{m}$ is needed in the extrinsic
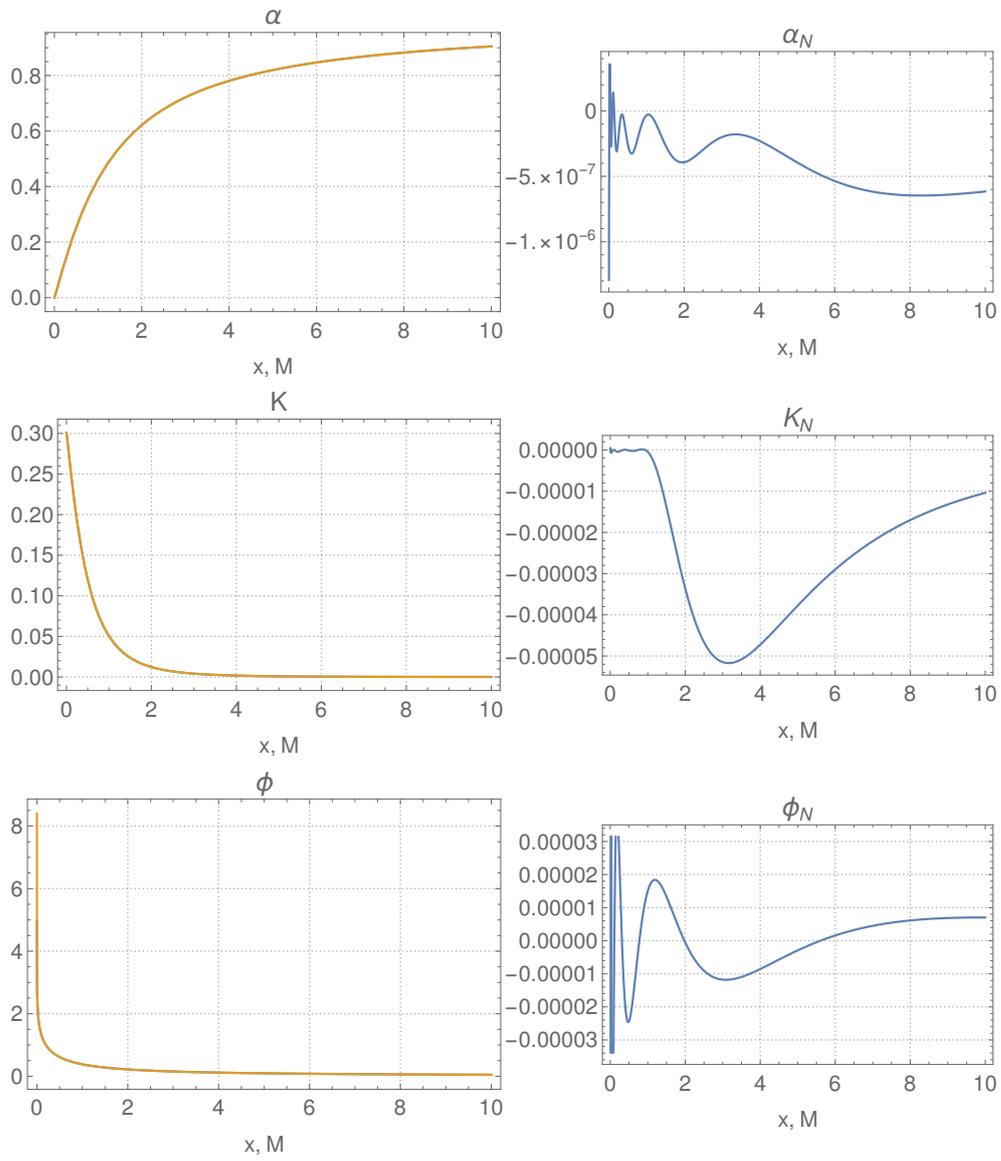
FIGURE 6.2   LEFT: *Functional form of the trumpet solution in blue, with the analytic pieces of 6·10 in yellow.* RIGHT: *Form of the numerical pieces in the case where grid spacing d → 0, computed by taking the difference between the analytic pieces given in 6·10 and the provided trumpet data. Continued in* FIG 6.3.
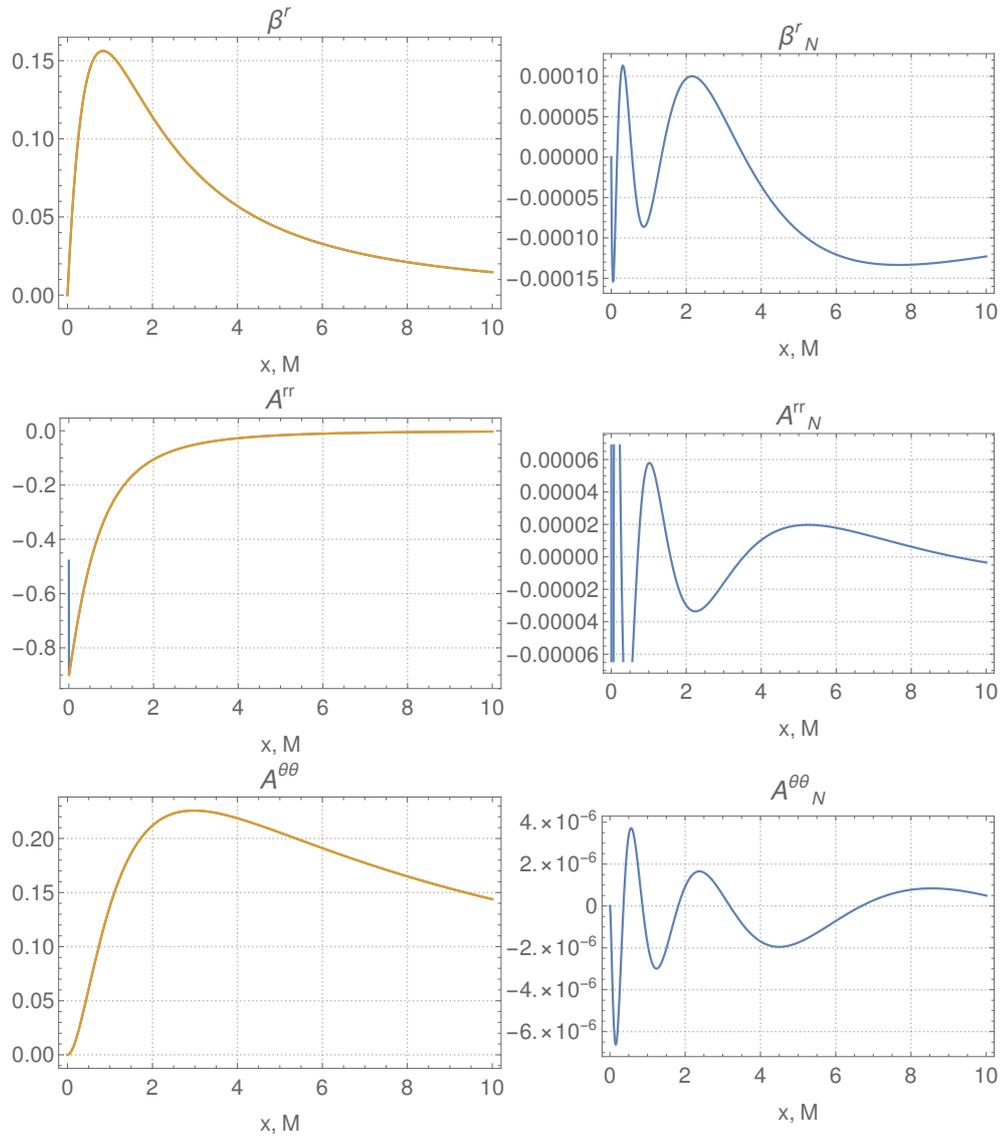
FIGURE 6.3   *Continued from* FIG 6.2. *The polar and azimuthal shift, as well as* $A^{\phi\phi}$, *are not shown as in this chapter we will be working with an entirely spherically-symmetric spacetime.*

curvature due to its nature as the Lie derivative of the metric along the normal to the spatial slice (1·2), so we also make the substitution $\mathbb{K} \to \frac{\mathbb{K}}{m}$, $\mathbb{A}_{ij} \to \frac{\mathbb{A}_{ij}}{m}$. A comparison between these newly-modified analytic pieces and the output from two static trumpet simulations of differing masses is shown in FIG 6.4 and FIG 6.5.

6.2.   *Static Schwarzschild*

The natural starting point for the testing of this method is the single static Schwarzschild black hole. This is known analytically in the 1+log trumpet case and so the numerical piece should be, numerical noise aside, identically zero. And as the solution is time-independent this numerical piece will be zero everywhere at all times. The purpose in starting with such a trivial case is simply to ensure that the modifications that have been made to the code work as expected and introduce no instabilities, allowing for early diagnosis and correction of any issue that may arise in this framework that will go on to be used in more complex scenarios in future. We have therefore $f \approx \mathbb{f}$, and so $f \approx 0$ with only inaccuracies in the fit and numerical noise contributing. As such it should be possible, with sufficiently good fits, to resolve the system with arbitrarily coarse resolution as *there are no features to be resolved numerically*.

Using standard methods we find that the black hole will evaporate in grid spacings $d \geq 4.7 \times 10^{-2} M$. Upon evolving with this new method we find remarkable stability in coarse resolutions— FIG 6.6 shows the lapse at various times throughout the evolution, both using these modifications and using standard methods, for a static Schwarzschild black hole in a grid spacing of $d = 32M$. These were run on 32 cores, with a single refinement level, time steps of duration $dt = 6.3 \times 10^{-2} M$, number of points $N = 44$, and $\eta = 0$ (see §2.5). An identical run with eleven refinement levels, and hence finest resolution $d = 3.1 \times 10^{-2} M$, is also shown for comparison. We see that use of the modifications to BAM introduced in this chapter together with a large grid spacing reproduces qualitatively the same behaviour as that found using unmodified BAM with fine grid spacing, preventing the evaporation of the puncture. To verify the performance of this we compare the results using standard methods, which we label 'old BAM', run at sufficient resolution for stability ($d = 4.7 \times 10^{-2} M$), and 'new BAM' run at $d = 24M$. The grid setup of these two runs is identical, save that the old BAM run has nine mesh refinement levels while the new BAM run has just one. The difference between their
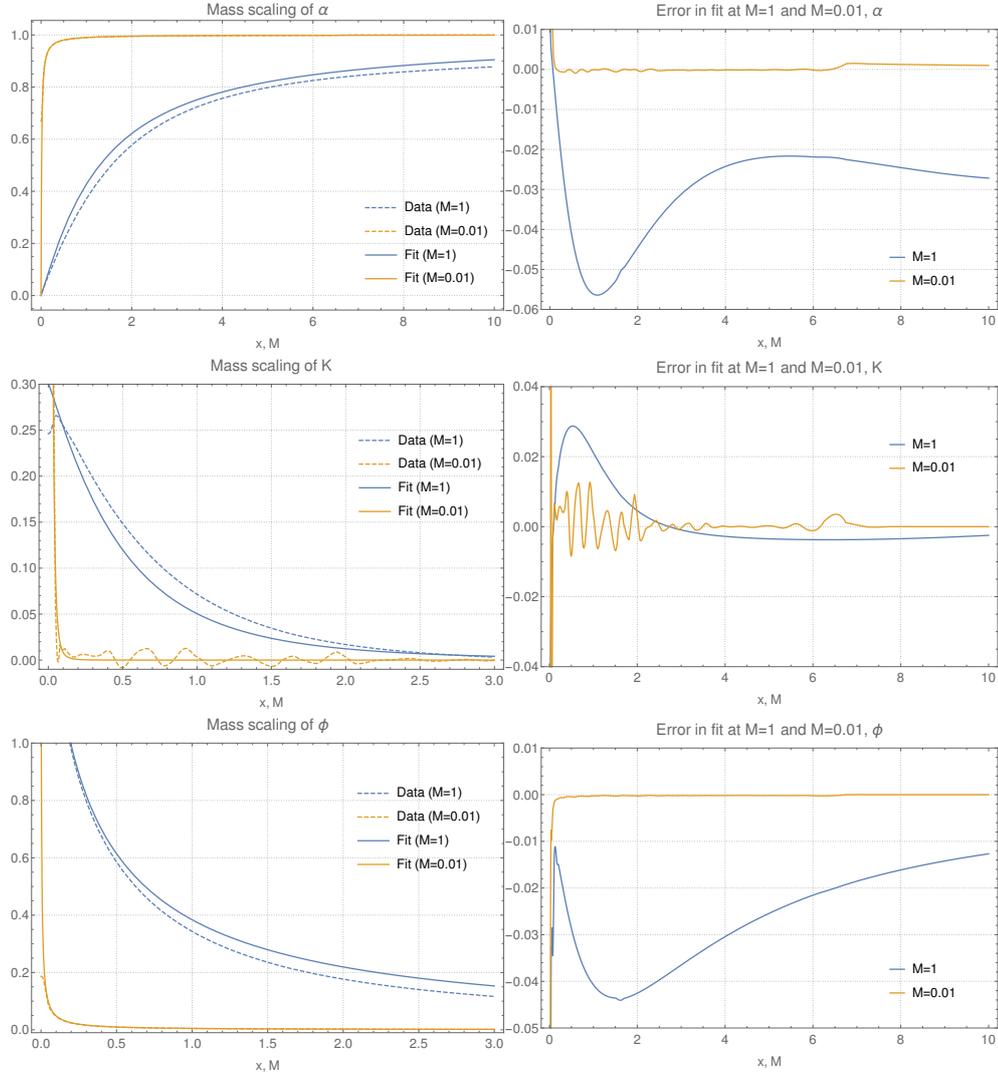
FIGURE 6.4   *Mass scaling of the BSSN variables in wormhole data at $t = 20M$ shown on the left, with the difference between data and analytic piece shown on the right. In blue are the results with $M = 1$ while yellow is used for $M = 0.01$. The dashed line indicates the analytic pieces of 6·10 with the mass-scaling modifications and the solid line indicates the data output from a static wormhole simulation. Note the kink in the lapse at $x = 6.8M$—this is the outwardly-propagating gauge error resulting from the use of wormhole data. Continued in* FIG 6.5*.*
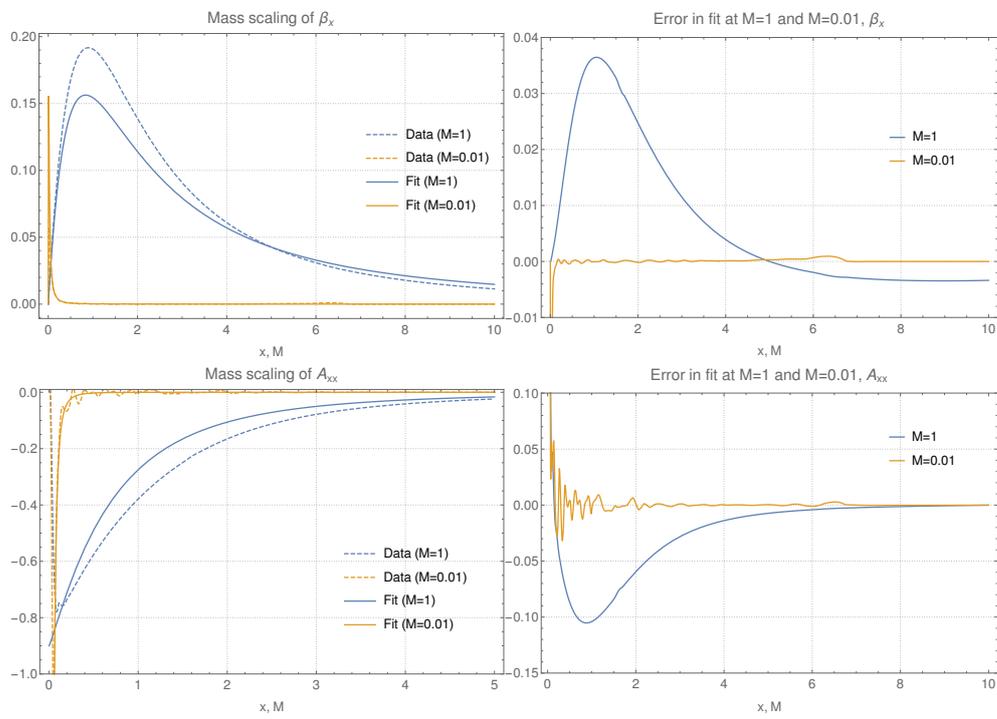
FIGURE 6.5 *Continued from* FIG 6.4.

respective BSSN variables is shown in FIG 6.7, using the data on the first refinement level in both cases.

These new methods demonstrate long-term stability (here defined as persisting, without signs of collapse or other instability, over at least $1000M$) for grid spacings within the range $2.5M \leqslant d \leqslant 40M$. Below approximately $2.5M$ the lapse steadily collapses to zero from the puncture outward, presumably the result of insufficient accuracy in the analytic pieces close to the puncture where gradients are steep and sensitivity to features in the numerical piece is high, and as such it is hoped that with more sophisticated fitting techniques the accuracy of the analytic pieces can be improved such that arbitrarily fine resolutions can once again be used. As the grid spacing is increased above $d \approx 20M$ it becomes necessary to severely restrict the Courant factor to the extent that a $d = 96M$ run was found to only be stable with $\Delta t = 2^{-16}M$. So while coarser resolutions may well be possible, the scaling of the required Courant factor makes them impractical, and for this reason we choose to state the soft upper limit as $d = 40M$, three orders of magnitude larger than previously possible with BAM.

These modifications to BAM are however currently incompatible with Berger-Oliger refinement of the time steps, and we are therefore limited to a uniform step duration. Exploratory tests on this suggest the issue lies in the communication of data between levels when they are not aligned with one another, presumably due to one level working in terms of the full BSSN evolution variable while those above or below are still working with the numerical piece. This is most likely a result of the way in which our modifications have been implemented: the analytic pieces are subtracted before evolution of the variables and then added back in afterwards, an order of operations that conflicts with Berger-Oliger time stepping where for each step at refinement level l the level $l + 1$ completes two steps.

These initial tests on this admittedly trivial case appear then to be encouraging. We have successfully evolved a black hole using resolutions three orders of magnitude coarser than previously possible, achieving similar levels of accuracy as traditional methods. It has been shown to display long-term stability; as a test of this the Schwarzschild case was evolved for a duration of $30,000M$ with $d = 24M$, running at an average wall-clock speed of $40,819.7M$ per hour on 32 cores. The same case run using old BAM with $d = 4.7 \times 10^{-2}M$ runs at an average speed of $67.7M$ per hour, representing a 60,000% speed increase. If this degree of improvement were to persist through to full Kerr inspirals, the
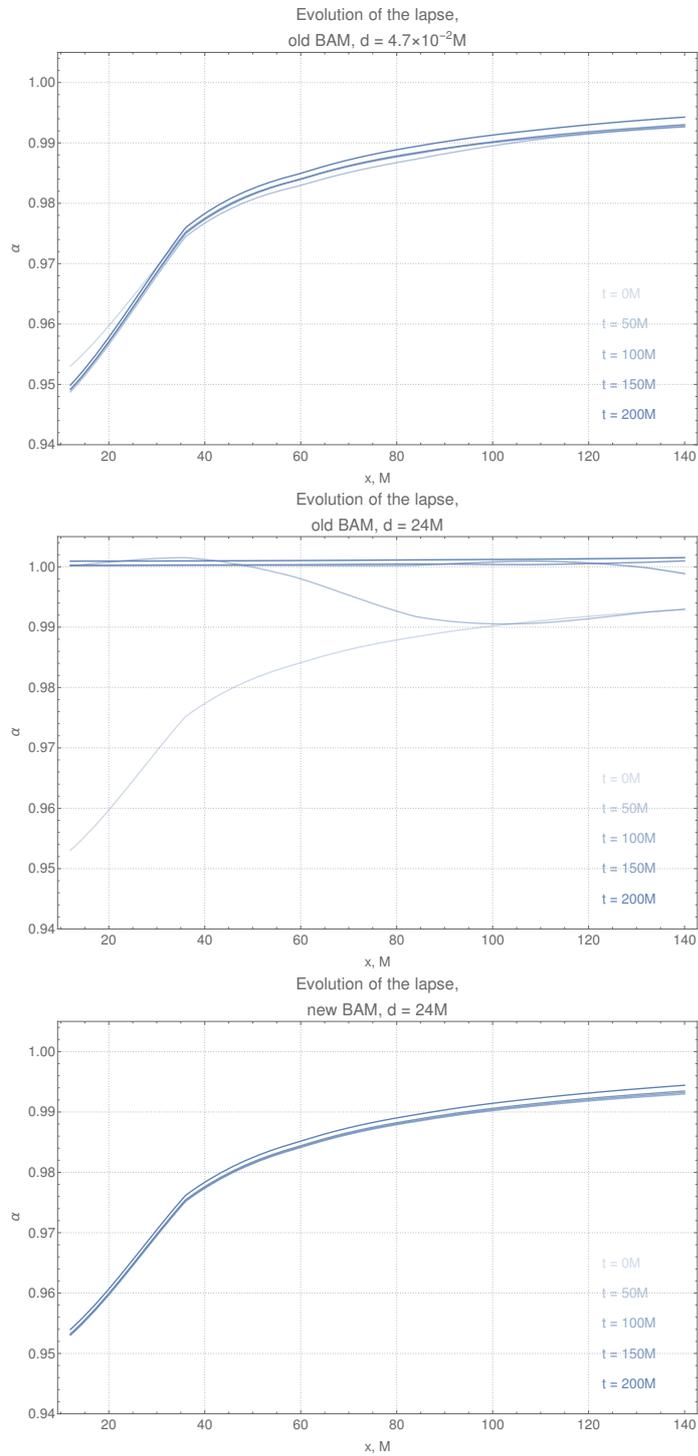
FIGURE 6.6 *Lapse at a selection of times throughout the first 200M of evolution of a static Schwarzschild puncture at a resolution of d = 24M, using these modifications (bottom panel) and using previous methods (centre panel), and for comparison the first 200M of a d = 4.7 × 10⁻²M 'old BAM' run (top panel). Opacity of the line represents progression in time.*
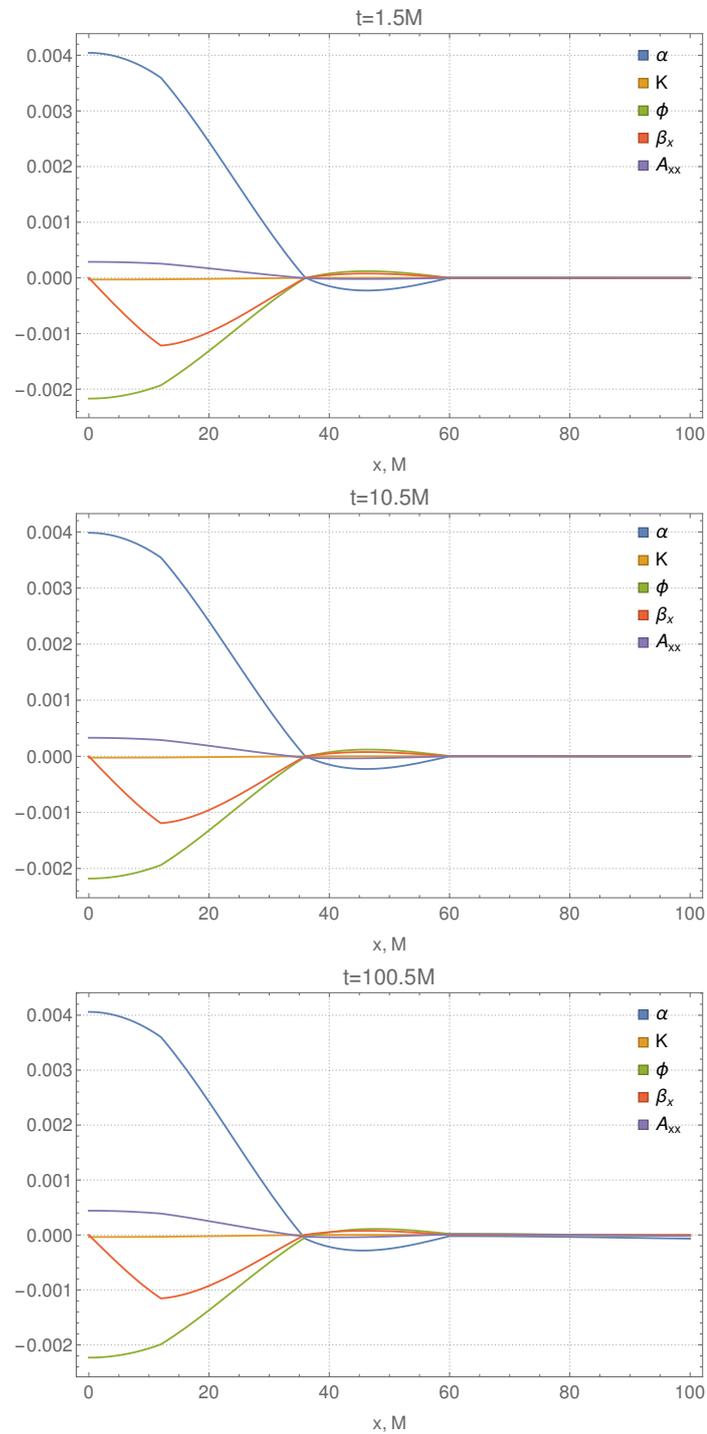
FIGURE 6.7 *Difference between old and new BAM runs at three different times during evolution. New BAM run with $d = 24M$, old BAM run at $d = 4.7 \times 10^{-2} M$ but data shown here taken from the level with $d = 24M$.*

$q = 8$ simulations in CHAPTER 4 that took approximately three months to complete would only take 3.5 hours. But a single static black hole is a long way from the ultimate goal of fully-inspiralling high mass ratio Kerr binaries; there remain significant obstacles to overcome before that will be feasible, chief among them the ability to evolve *moving* black holes across the domain. This in itself represents a set of not-insignificant challenges that we introduce and begin tackling in the following chapter.

Extending to Dynamic Black Hole Spacetimes

7

With the success of the stationary case the next logical step is a single boosted Schwarzschild black hole. Superficially this seems trivial—there's still only one black hole, it's still non-spinning, there is no acceleration and motion is constrained to one dimension—but in practice there are a number of subtleties both in the implementation within BAM and the construction of the analytic pieces that could previously be neglected. In this chapter we introduce and address each of these, performing the same kind of preliminary tests as those done for the static case, before extending this method to the head-on collision of two black holes of mass ratio $q = 100$.

The aim in this chapter is not to produce highly accurate simulations on par with those produced using standard methods, but rather to show that in principle the approach introduced in the previous chapter can work, that it can produce stable simulations of dynamic spacetimes that approximate the correct phenomenology at a fraction of the computational cost. This is demonstrated in tests on a boosted Schwarzschild trumpet with a velocity 1% the speed of light, stable over $3500M$ of evolution, and for the head-on collision of two black holes, both in resolutions comparable to those achieved in the static case, $d = \mathcal{O}\left(10M\right)$. The inability to evolve traditionally fine resolutions encountered in the previous chapter remains, as does the incompatibility with Berger-Oliger time-stepping, but the largest caveat to this is that this preliminary success depends entirely on manually prescribing the puncture trajectories using a previously-run high resolution version of the simulation. New BAM as it currently exists is unable to capture the location of the puncture and as such is unsuitable for its intended purpose of making practical those simulations which are currently impractical. Further, we find that the numerical

pieces at the puncture in the current implementation of this method are not flat, simple functions that can be resolved in coarse resolutions, but are sufficiently flat for stability only. The importance of the results presented in this chapter lies rather in the demonstration that it is fundamentally possible to stably evolve dynamic black hole spacetimes in extremely coarse resolutions *given the puncture can be accurately tracked in some way*. This is a problem left for the most part to future work. Prescribing the trajectory allows us to investigate the viability of our method—namely, that the black hole doesn't evaporate, that it moves as it should, that there are no disruptive perturbations—without first having to develop a novel tracking routine that may prove entirely unneeded. We propose one potential coarse puncture tracker that currently fails when applied to anything beyond a static trumpet located at the origin and suggest avenues for further development of the technique in the hope that it can be refined.

7.1. *Boosted Schwarzschild*

The first and simplest modification that must be made to the framework outlined in CHAPTER 6 is to allow for the puncture to be located somewhere other than the origin. This requires a simple redefinition of the radial coordinate r in our analytic pieces taking into account puncture location $\{x_1, y_1, z_1\}$,

$$ r = \sqrt{x^2 + y^2 + z^2} \rightarrow \sqrt{(x - x_1)^2 + (y - y_1)^2 + (z - z_1)^2} \, , \qquad 7.1 $$

though for all simulations considered in this chapter the punctures will be confined to the x-axis, with $y_1 = z_1 = 0$. This however immediately introduces a problem. Where does this location come from? Certainly the puncture is tracked and output by BAM; this is straightforward in current moving-puncture codes as the velocity can be directly read off as the value of the shift vector at the puncture, making computation of the position at each time step trivial. This is discussed in further detail in §2.5. But the philosophy behind this new approach is specifically to *not* resolve the spacetime near the puncture, so we no longer have access to an accurate reading of the shift at the puncture. An entirely new approach is required here, an example of which is introduced in §7.3. For the moment though, while we are conducting preliminary tests of the method on simple configurations, we settle for a straightforward *ad hoc* solution that allows us to focus on

setting up the bulk of the new framework and ensuring that this new approach can be successfully applied to moving punctures. Specifically, for each attempted new BAM run we first simulate the system using old BAM with traditionally fine resolutions, achieved by adding additional refinement levels with no further alterations, then make an analytic fit to the puncture positions and velocities this outputs, $\mathbb{X}(t)$ and $\mathbb{V}(t)$.

One further factor that didn't pose an issue in the static case is the fixed nature of the grid points in BAM, as by necessity a boosted puncture will at times be closer to a grid point and at times further away. This means that our analytic pieces must now be accurate arbitrarily close to the puncture, as any significant deviations at any point along the fit will be seen by a grid point. The fits outlined in 6·10 have already been optimised in this regard, though further refinement appears to be needed given the observed instability of new BAM simulations with small grid spacings. Passing directly over a grid point potentially risks destabilising the evolution; although in the moving puncture scheme the singularity doesn't pose such a risk (see §2.3), we're now manually inputting functions that in some cases diverge at $x = x_1$. This will need to be handled in future, but BAM's staggered grid setup ensures that none of the cases considered in this chapter face this issue as they are all constrained to the x-axis. This threads the centre of each cell, so a puncture boosted along it moves through a tube clear of grid points. But for generic motion and full inspirals this will need to be addressed, even if the coarseness of the grid means that a direct interaction with a grid point is unlikely, perhaps through the use of fits that are well-behaved at the puncture.

There is however one alteration that needs to be made to the fits presented in 6·10. A boost along the x-direction, which we restrict our focus to in this chapter, will require the special-relativistic scaling of the x-coordinate in our fits by the Lorentz factor, $x \rightarrow \gamma x$ where

$$\gamma = \frac{1}{\sqrt{1 - v^2}} \, ,$$

7·2

and just as with the mass scaling in CHAPTER 6 the derivatives along the x-direction contribute a further factor of γ. Including this is not just a matter of completeness—the inspiralling black hole binary simulations used within the LVK regularly achieve linear velocities of $0.2c$ [79], so γ can contribute considerably to the overall shape of the BSSN variables.
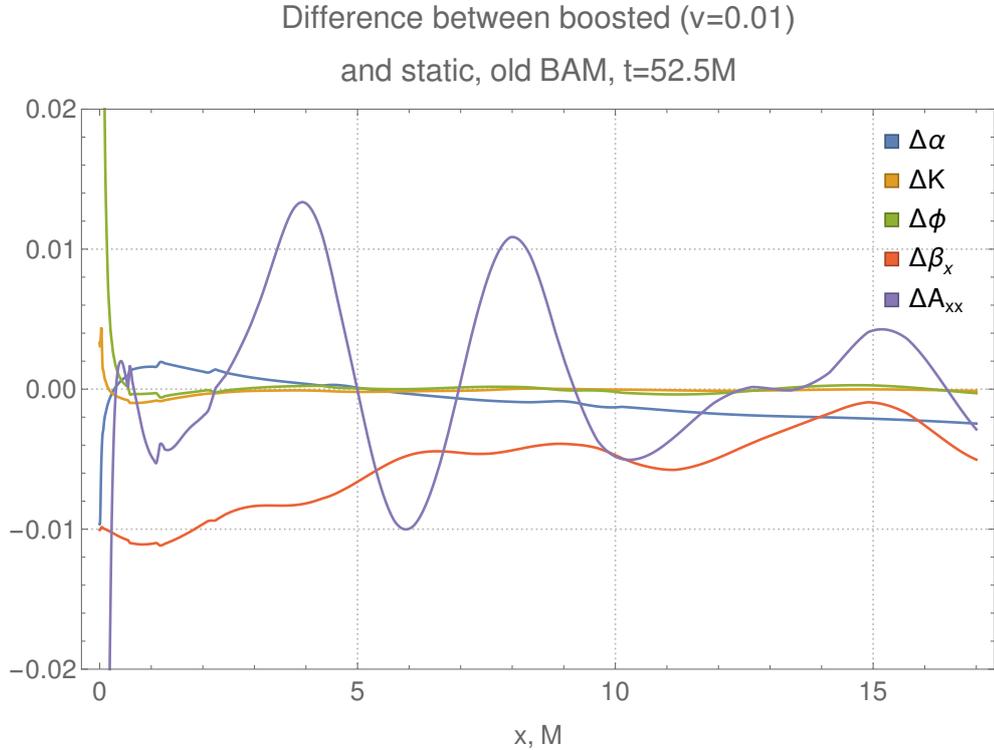
FIGURE 7.1   *Difference between BSSN variables output by a static Schwarzschild simulation and a boosted ($v = 0.01$) black hole simulation. These snapshots are provided at $t = 52.5M$ to allow time for the gauge to relax and for junk radiation to dissipate. Only β_x and A_{xx} are shown here as the boost is applied in the x-direction, preserving the symmetry in the y- and z-directions.*

Aside from these relatively minor alterations we leave the majority of the fits from the previous chapter, which were based on the static solution, unchanged despite the boost, anticipating that the changes it introduces will be sufficiently slowly-varying and of sufficiently small magnitude that the numerical pieces remain resolvable with large grid spacings. If we examine the differences between the static and boosted cases we see that this is justified in all cases save the shift: FIG 7.1 shows our evolution variables in the static and boosted cases, along with the difference between them.

It is clear that a fit to the shift vector based solely upon the stationary data leaves unacceptably sharp features in the numerical piece. There is also the suggestion of a need for modification of the lapse, though at this preliminary stage this is neglected and to future developments if necessary. A more thorough treatment of the changes in the
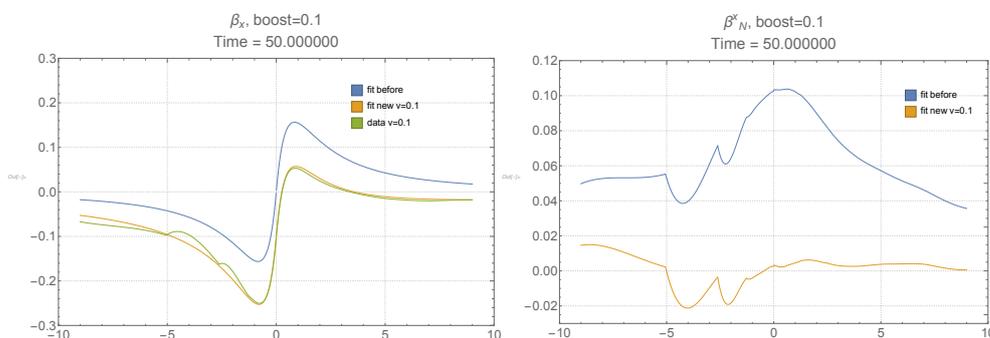
FIGURE 7.2   LEFT: *static analytic shift in blue, boost-modified analytic shift with v = 0.1 in yellow, and β_x from a boosted (v =0.1) simulation run using old BAM. This velocity was chosen so as to more clearly demonstrate the boost-induced features while still remaining astrophysically reasonable based on linear velocities achieved during comparable-mass binary evolutions (see for example [79]). RIGHT: The β_x resulting from use of the static analytic shift (blue) and the boost-modified analytic shift (yellow).*

shift due to the boost is needed, so we appropriate from [131] an expression for the shift of a wormhole boosted with velocity v along the x-direction,

$$\beta_{boost} = \frac{1 - \frac{\alpha^2}{\psi^4} + 2\beta^2 + \frac{\beta}{v} - \beta v}{2\beta + \frac{1}{v} - \frac{\alpha^2}{\psi^4}v + \beta^2 v} \ ,$$

7.3

and assume that trumpets, on which our analytic pieces are based, behave similarly. We apply this to our fit from the static case to form our boosted β_x, with the comparison against the boosted data shown in FIG 7.2. This displays good agreement with the data with the problematic features removed from the numerical piece.

If we now attempt a coarse-grained evolution we see that it is stable for approximately $3000M$. FIG 7.4 shows the lapse at various times throughout the evolution, using old and new BAM, for a boosted ($v = 0.01$) Schwarzschild trumpet in a grid spacing of $d = 24M$. These were run on 32 cores, with five refinement levels, time steps of duration $dt = 1.1 \times 10^{-2}M$, number of points $N = 40$, and $\eta = 0$. An identical run with 15 refinement levels ($d = 2.3 \times 10^{-2}M$), is also shown for comparison. It is from this finely-resolved old BAM run that we take the puncture locations and velocities at each time step to construct the proxy tracker that is used in the coarse simulations. The fits to this data
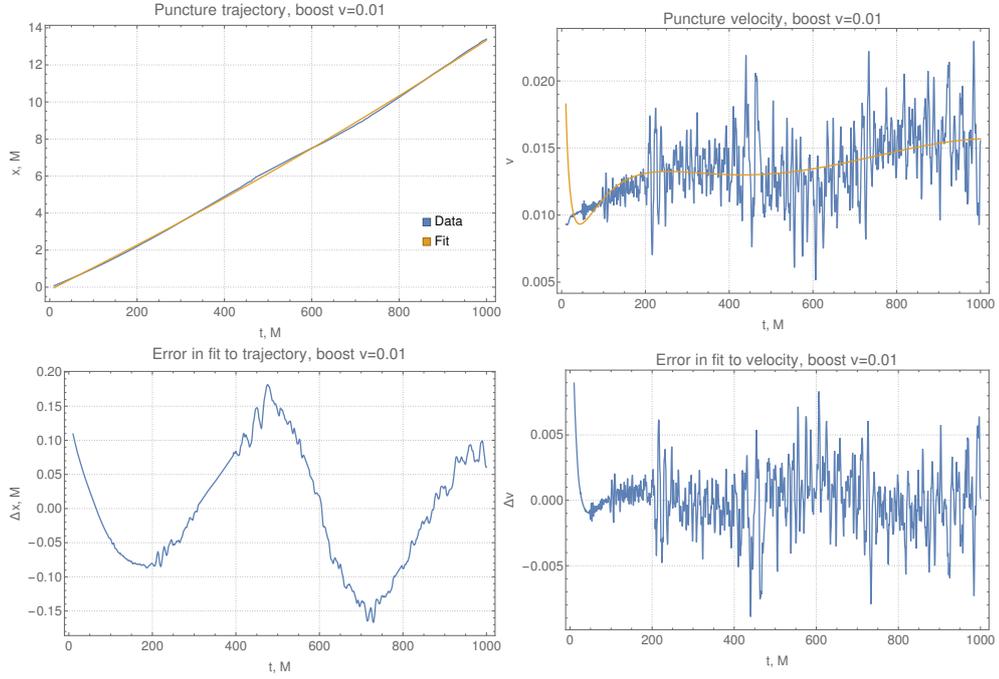
FIGURE 7.3   *Form of the proxy tracker used to evolve the v = 0.01 boosted trumpet. On the left is shown the fit through the position data with its deviation from the data below, and on the right is the corresponding fit to the velocity, given in 7·4 and 7·5 respectively.*

that are used are as follows, and are shown in FIG 7.3.

$$
\begin{aligned}
\mathbb{x}(t) = &-0.1452 + 0.0205t^{1.356} - 0.0265t^{1.287} - 0.0186t^{1.285} - 0.0165t^{1.285} \\
&- 0.0159t^{1.285} - 0.0165t^{1.284} - 0.0027t^{1.279} + 0.0098t^{1.258} + 0.0098t^{1.258} \\
&+ 0.0219t^{1.247} + 0.0448t^{1.236} + 0.0052t^{-2.177} ,
\end{aligned}
$$

    7·4

$$
\begin{aligned}
\mathbb{v}(t) = &\ 0.3676t^{1.132} - 0.4756t^{1.105} - 0.3862t^{1.079} + 1.5017t^{0.992} - 0.0753t^{0.963} \\
&- 0.0885t^{0.958} - 0.2926t^{0.908} - 0.2947t^{0.907} - 0.3323t^{0.897} - 0.3321t^{0.896} \\
&- 0.2713t^{0.880} + 0.7122t^{0.806} - 0.0834 .
\end{aligned}
$$

    7·5

FIG 7.4 shows the lapse at various times throughout the evolution, using these modifications to the methods described in §6.2, for a boosted ($v = 0.01$) Schwarzschild black hole in grid spacing of $d = 24M$. The puncture no longer evaporates and does move across the grid as expected, though there are clear qualitative differences when compared
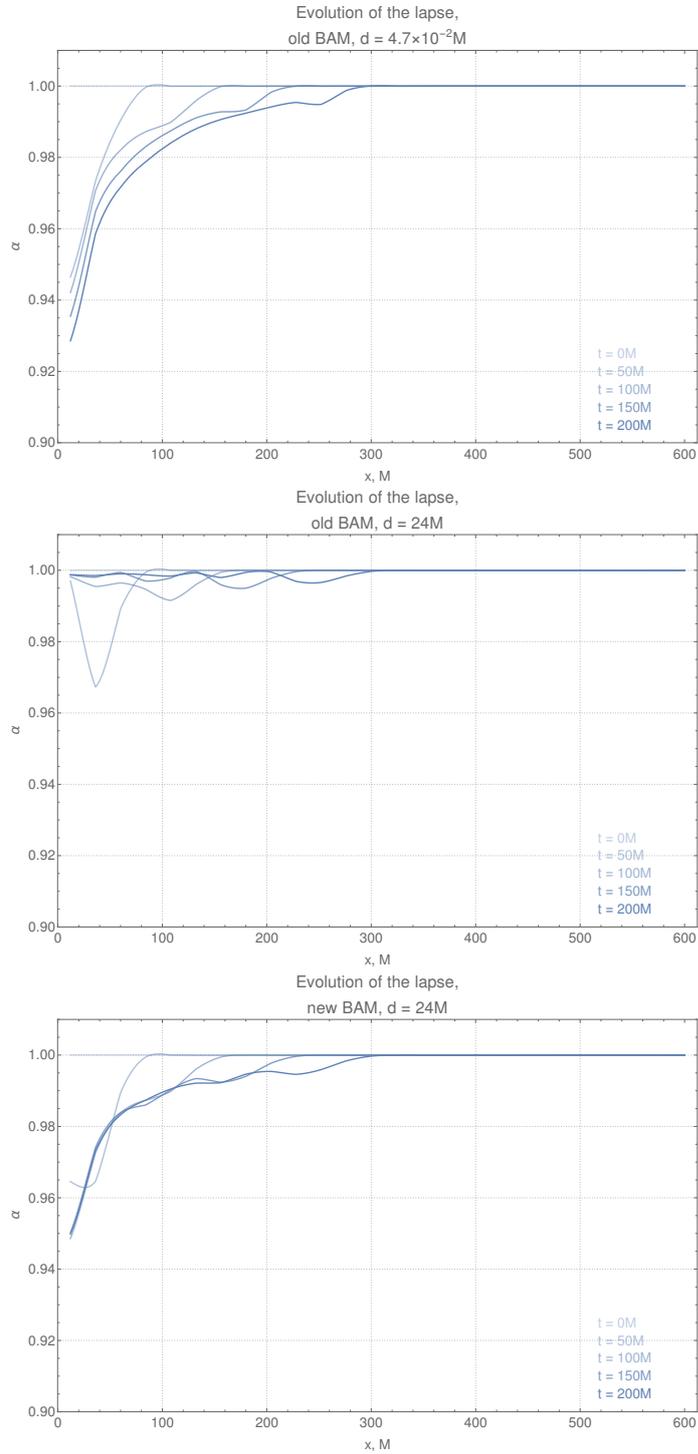
FIGURE 7.4 *Lapse at a selection of times throughout the first 200M of evolution of a boosted (v=0.01) black hole at a resolution of d = 24M, using these modifications (bottom panel) and using previous methods (centre panel), and for comparison the first 1000M of a d = 4.7 × 10^{-2}M old BAM run (top panel). Opacity of the line represents progression in time.*

to the finely-resolved old BAM run: as the puncture moves towards the nearest grid point the lapse appears to deepen, which is not observed in the new BAM run. This is most likely due to the presence of the finer levels in the old BAM simulation, which resolve the steep gradients close to the puncture and communicate this data back to the coarser levels resulting in coarse-level data informed by the behaviour of the finer levels. The instability at fine resolutions that was observed in the static case persists here and as before, for verification we also compare a coarse new BAM run with old BAM run at sufficient resolution for stability, $d = 4.7 \times 10^{-2} M$. The grid setups of these are identical save that the old BAM run employs twelve refinement levels rather than the three used by the new BAM run. The difference between their respective BSSN variables is shown in FIG 7.5, using the $l = 3$ data in both cases.

Having established that new technique is fundamentally successful in both the static and boosted cases, there is nothing to suggest *a priori* that it can't be applied to the more complex inspiralling binary configurations. The ability to update the numerical pieces at each time step and retain stability as this is done, together with the observation that the introduction of a boost does not necessarily lead to short-wavelength features in the numerical pieces, represents a significant development from the implementation of this method in CHAPTER 6. Both of these aspects were potentially fundamental obstacles to this approach; the static case is independent of time and known analytically in the 1+log gauge, but it is in attempting to apply these methods to dynamic spacetimes that we can test if it remains robust under perturbation. With the ability to evolve a moving puncture we can take the first step towards the binary inspiral simulations we ultimately aim to achieve.

7.2. *Head-on Collision*

In this section we use the boosted techniques we have constructed above to simulate the head-on collision between two black holes, taking two static punctures and allowing them to accelerate towards one another. This is of course not an astrophysically realistic configuration, but it represents the next logical step from the boosted case and, importantly, a step towards the ultimate goal of a binary inspiral. An orbiting binary introduces a host of complications that would only serve to muddy the water at this stage, obfuscating the source of errors when they inevitably appear and slowing development. Tracking would
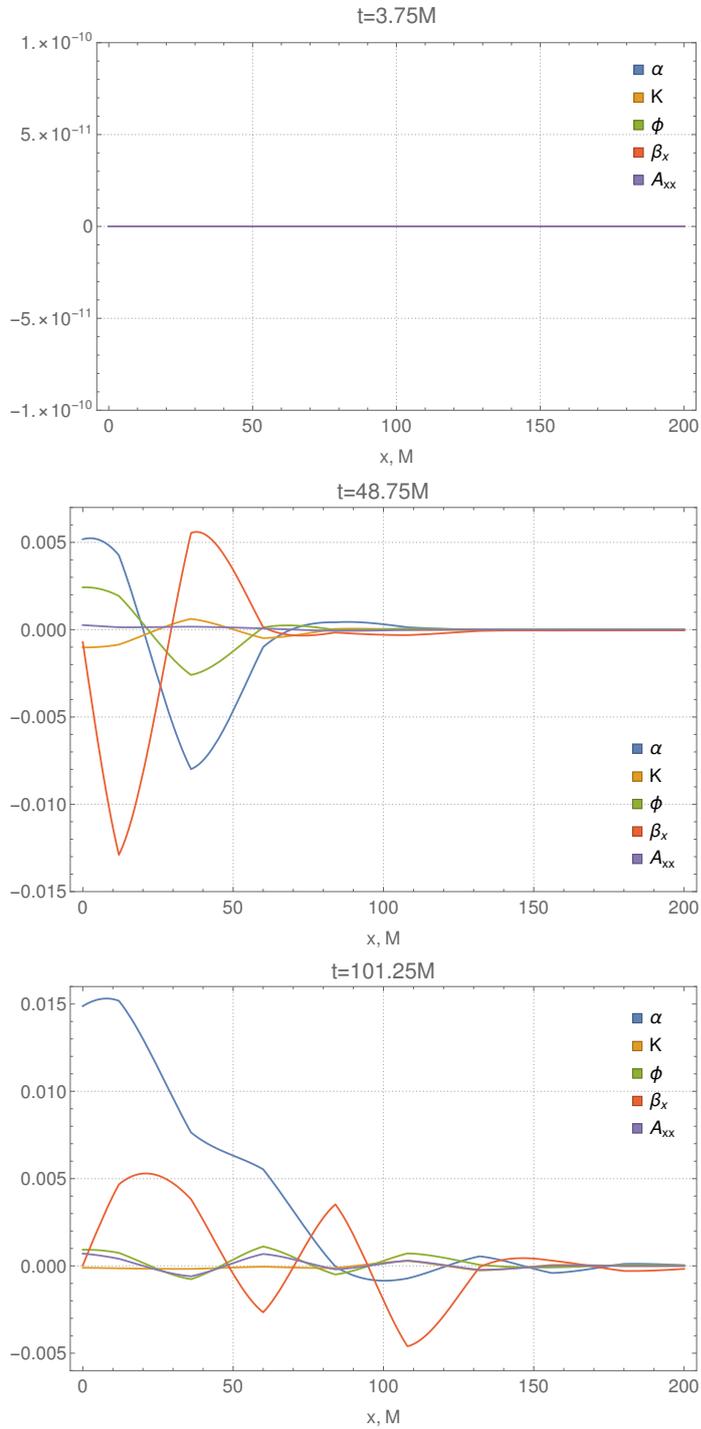
FIGURE 7.5 *Difference between old and new BAM runs of the $v = 0.01$ boosted puncture at three different times during evolution. New BAM run with $d = 24M$, old BAM with $d = 4.7 \times 10^{-2} M$ but data shown here from the level with $d = 24M$.*

be a particular issue given we still rely on standard fine-resolution simulations to generate the proxy tracker, which for binaries would require timescales impractical for exploratory development work. The motion of the puncture would no longer be constrained to a tube free of grid points as it was in the previous section, risking a direct encounter with a grid point (though the probability of this will be reduced due to the extremely coarse grids we hope to employ). We would also no longer be able to take advantage of any symmetries in the system nor the initial stationarity of the punctures. The head-on collision removes a great deal of these complications while still allowing us to test whether the presence of a second puncture introduces disruptive instabilities, for example through the presence of interaction terms in the solution that are not currently captured by our analytic pieces which are based on the single static Schwarzschild trumpet. We present the results of tests on the collision of a $q = 100$ pair of black holes, found to demonstrate stability through to merger in similarly coarse resolutions as the static and boosted tests, with no disruption as they approach one another.

With the addition of a second compact object further along the x-axis an initially static black hole will accelerate inwards, travelling along the axis until the two objects merge. This is in many ways similar to the boosted case we've discussed, except that the velocity is no longer constant and the background metric will now feature the presence of the other black hole. This could pose a problem for our approach. Up until now we have relied on the background metric being flat and featureless once the foreground black hole is removed making it possible to successfully evolve in coarse resolutions. Even if the second black hole is evolved with traditionally fine resolutions, the background it induces around the primary black hole where we are applying our method could be too finely-featured to be resolved with these coarse grid spacings. One solution would be to increase the separation to such a degree that the background is sufficiently smooth. But eventually as this is evolved the puncture will encounter the same sharp features we were trying to avoid, so instead we increase the mass of the second black hole. The length scale of the background metric is determined by the mass of the secondary black hole, so by increasing the mass ratio we can effectively smooth the background experienced by the primary black hole. At this exploratory stage of development we choose to work with a $q = 100$ binary with a separation between the components of $10M$, primarily because we still need to run this configuration with old BAM to generate the proxy tracker and so are limited to cases that can run to completion using old BAM in practical timescales.

This was run at a wall-clock speed of $0.6M$ per hour on 32 cores, with resolution around the larger black hole of $4.6 \times 10^{-2}M = 4.7 \times 10^{-2}m_1$ and around the smaller black hole $3.6 \times 10^{-4}M = 3.7 \times 10^{-2}m_2$. The domain consists of 9 refinement levels around the large black hole and 16 around the smaller. The fits to position and velocity that form our proxy tracker are

$$\mathbb{X}(t) = 10 - \frac{0.0032t^{2.206} + 0.0031t^{2.165}}{1.8102t^{0.176} + 6.1308t^{-0.863}} \, , \qquad 7 \cdot 6$$

$$\mathbb{V}(t) = \frac{-0.1455t^{1.334} - 0.1365t^{1.333} + 0.67917t^{1.327} - 0.6572t^{1.277} + 0.1659t^{1.121} + 0.2054t^{1.073}}{0.0001t^{3.145} - 0.0002t^{2.872} - 0.0003t^{2.823} + 0.0001t^{2.707} + 18.6319t^{-0.702} + 18.6319t^{-0.702}} \cdot$$
$$7 \cdot 7$$

When running with new BAM we do precisely the same but drop down to 9 levels around both black holes—for the small black hole, where we will be applying the analytics, this corresponds to an effective resolution of $4.7m_2$ (cf. CHAPTER 6 where resolutions in the range $1.5M \lesssim d \lesssim 48M$ were successfully employed).

Unfortunately we are immediately presented with a problem. Our analytic pieces are based on trumpets, but there exists at present no initial data for binary trumpets. Until such a time as this is developed we must fall back on wormhole data instead, which evolve to trumpet form after some initial relaxation time (§2.3). But this means we need to wait for this transition to have completed before our low-resolution methods can be applied, or else our analytic pieces won't accurately represent the full solution and the numerical piece will no longer be sufficiently flat and featureless. What is needed then is the ability to evolve using old BAM while the puncture is settling down to trumpet form—with traditional, fine resolutions and no analytic pieces, as these are currently incompatible with each other—before dropping down to a much lower resolution and switching on our modified evolution variables. Clearly this is not optimal and represents a severe bottleneck in the generation of high mass ratio simulations, particularly once this approach is pushed to EMRI-scale mass ratios, but for proof-of-principle tests it will be sufficient. Once binary trumpet initial data is developed we will be in a position to swap out the wormhole data we're currently using and take advantage of a framework that has already been developed and tested.

This shift to lower resolution is achieved by shedding refinement levels at a user-specified time. We modify an existing function within BAM intended to be used for the shedding of levels after merger, when the remnant black hole would be larger than

the parent objects and so wouldn't require such fine grid spacings to be resolved. At some given time[1] this function deallocates the memory used by the variables associated with those levels below the user-specified new finest level, $l'_{max} < l_{max}$. Some care is needed here however: with each finer level being advanced one time step successively, and Berger-Oliger levels being advanced by a greater number of smaller time steps, shedding levels arbitrarily can lead to the inconsistent communication of data between levels. It is essential to wait for the next instant after the specified shedding time that all levels are aligned, that the finest level and coarsest level are both at the same time step, before levels are shed. There is also an extreme sensitivity to the order of events involved. We've discussed how the new BAM modifications are currently incompatible with Berger-Oliger time stepping, and how the analytic pieces used aren't accurate enough to be used at fine resolutions, so these modifications must switch on *after* the levels have been shed, when we are now operating in coarse resolutions and on levels that employ uniform time stepping. Equally the solution can't be resolved at l'_{max} without the modifications, so there can be no evolution steps between the shedding of levels and the activation of the analytics. This, together with the requirement that all the levels be aligned, means that before the new BAM modifications can be switched on both the specified 'merger' time must have been reached and the shedding must have taken place: $l_{max} = l'_{max}$.

In attempting to evolve this head-on $q = 100$ collision we see that it is stable through to merger. FIG 7.6 shows the lapse at various times throughout the evolution, using old and new BAM, for a pair of initially-stationary wormholes of mass ratio $q = 100$ separated by $10M$. These were run on 32 cores, with eight refinement levels covering the larger puncture ($d = 0.24M$) and sixteen over the smaller puncture ($d = 9.4 \times 10^{-4}M = 9.4 \times 10^{-2}m_1$). After a duration of $0.3M = 30m_1$ we shed eight of the levels covering the smaller puncture, taking the finest resolution to $d = 0.24M = 24.2m_1$. Time steps of duration $dt = 4.7 \times 10^{-4}M$, $N = 48$ points, and $\eta = 0$ were used. An identical old BAM run without shedding is also shown for comparison. Remarkable agreement is seen in the phenomenology of the simulation as the smaller black hole appears to accelerate, approach the larger black hole, and merge without crashing and without the introduction of instabilities. Again, puncture locations and velocities at each time step are taken from the finely-resolved old BAM run to construct the proxy tracker below, shown in FIG 7.7,

---

[1]For binary simulations it is possible to instead specify the separation between the two punctures at which shedding should occur.

and the difference between the BSSN variables in each case is shown in FIG 7.8, using the $l = 8$ data in both cases.

The numerical pieces that would result from subtracting our analytic pieces from the output of a simulation of a head-on collision of two ($q = 100$) wormholes separated by $10M$ are shown in FIG 7.9 for a run with finest-level resolution of $d = 0.24M$ covering the larger puncture and $d = 9.4 \times 10^{-4}M$ covering the smaller. The extrinsic curvature looks to be particularly problematic, primarily due to its mass scaling—at low masses the additional factor of $\frac{1}{m}$ we found to be required in CHAPTER 6 causes leads to a strong narrowing of the characteristic width of the function and strong inflation of its magnitude, so even a small misalignment in the position of the analytic piece and the 'true' position of the puncture will disrupt the numerical piece. The numerical shift is initially poor because in using wormhole data we have an initially vanishing shift[2] that evolves into its expected form, an issue circumvented through the use of level-shedding. All numerical pieces remain relatively featureless beyond a region surrounding the puncture approximately $0.1M$ in diameter, but this is nonetheless outside the horizon of the small black hole. This suggests that the apparently successful simulation of the head-on collision using new BAM may in fact *not* represent the true solution accurately. Qualitatively it behaves as expected, due in no small part to the prescribed puncture locations it relies on, but also because despite the features in the numerical piece existing outside the horizon, they exist entirely between two grid points in our $d = 0.2M$ simulation shown in FIG 7.8. At the two nearest grid points to the puncture the numerical pieces will indeed be as flat and slowly varying as hoped, and so the simulation can be evolved stably—the features we see in the numerical piece cannot therefore be disruptive, but the question remains as to whether they prevent sufficiently accurate evolutions and, as a result, prevent the construction of an independent puncture tracker.

All three of the scenarios that have been tested with new BAM have displayed long-term stability and approximately the same phenomenology as their equivalent old BAM runs. This is encouraging—clearly the use of analytic knowledge of the BSSN variables and their derivatives to reduce the resolution required by black hole simulations is not a fundamentally flawed endeavour. It may yet prove that analytic pieces cannot be constructed for full binary inspirals that are of sufficient accuracy for gravitational wave astronomy,

---

[2]This is done for the sake of simplicity, but really any choice is equally valid here as the gauge conditions act to push the shift to its preferred value.
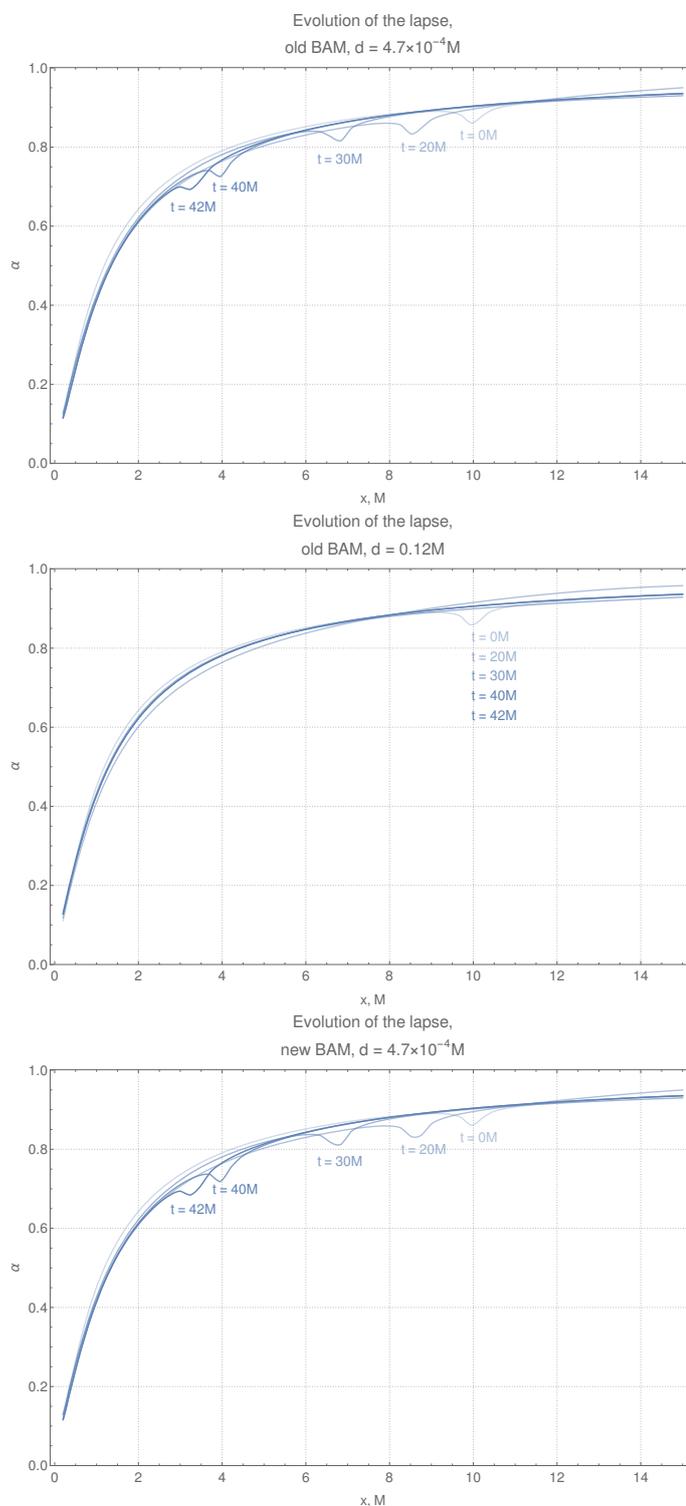
FIGURE 7.6   *The lapse at a selection of times throughout the evolution of the head-on collision, which merges at approximately $50M$, with a resolution on the finest level of $0.12M$. The bottom panel shows the results when using new BAM while the centre panel shows the same configuration run using old BAM. In both cases we begin with 8 levels on the larger black hole and 16 on the smaller before shedding 8 of the levels around the smaller black hole at $t = 0.2M = 20.2m_1$. For comparison a $d = 4.7 \times 10^{-4} M$ old BAM run is shown in the top panel. Opacity of the line represents progression in time. The smaller black hole evaporates immediately after this shedding of levels when running with old BAM but persists through to merger when running with new BAM.*
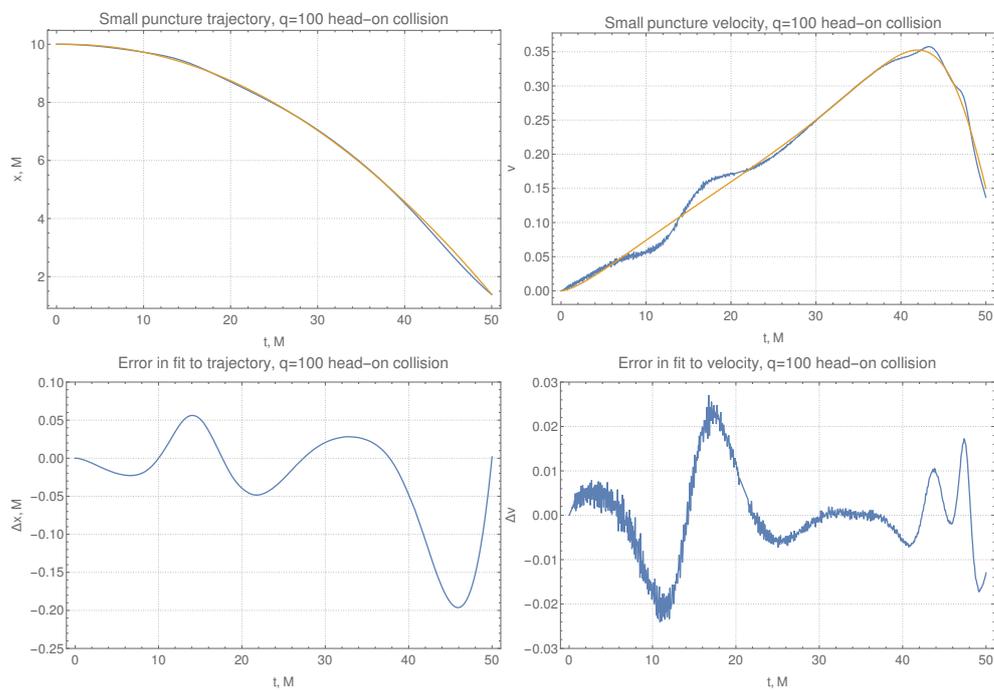
FIGURE 7.7 *Form of the proxy tracker used to evolve the head-on collision. On the left is shown the fit through the position data with its deviation from the data below, and on the right is the corresponding fit to the velocity, given in 7·6 and 7·7 respectively.*
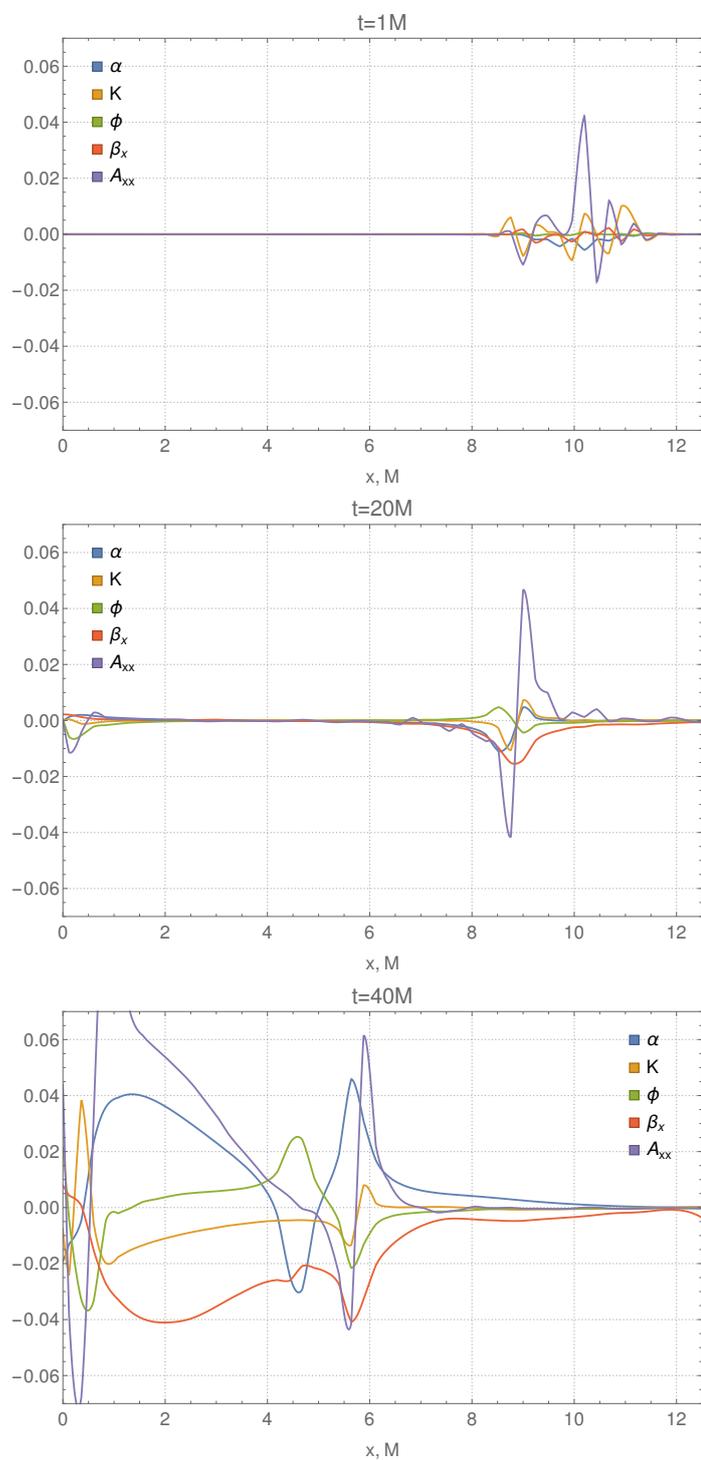
FIGURE 7.8   *Difference between old and new BAM runs at three different times during evolution for the head-on collision. The new BAM run used a resolution on the finest level of $23.8m_1$ and the old BAM run used $9.3 \times 10^{-2} m_1$, though the data shown here is that taken from the level with $d = 23.8m_1$.*
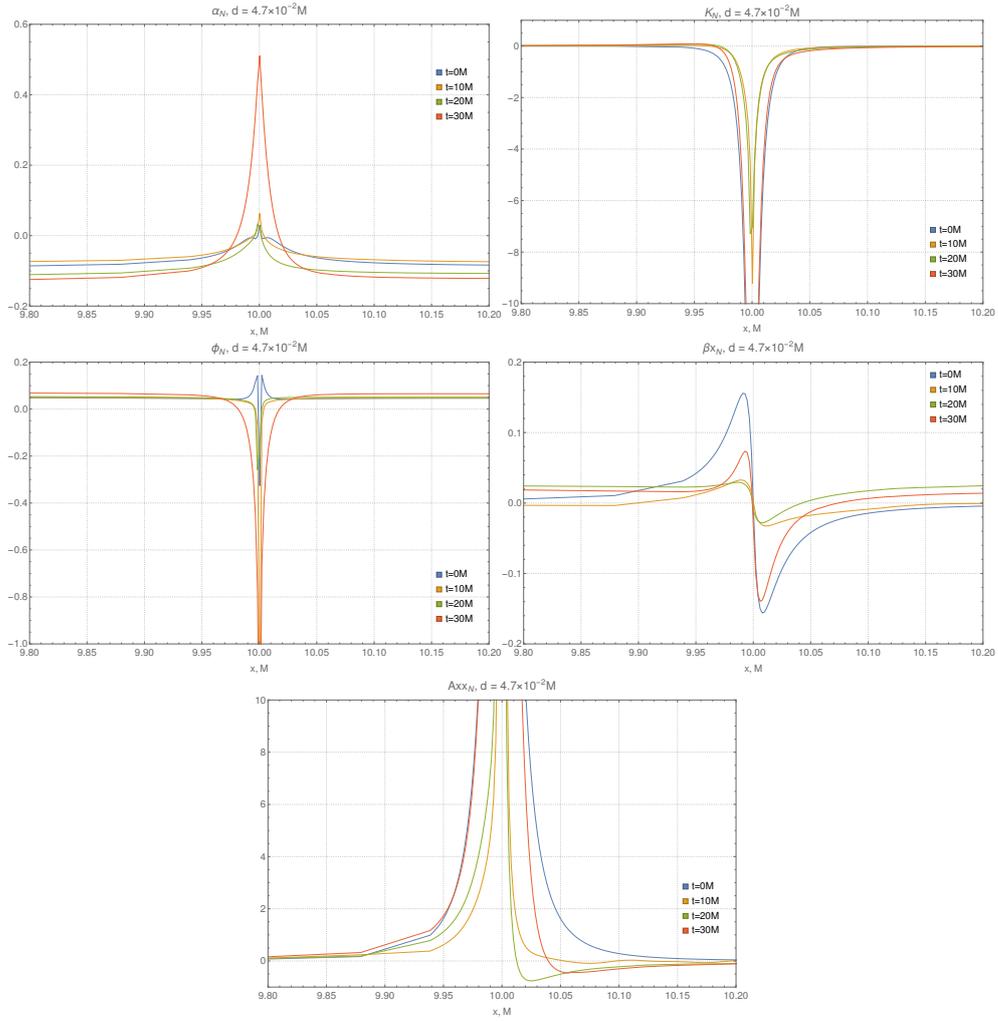
FIGURE 7.9 *Evolution of the numerical pieces that would result from a $d = 4.7 \times 10^{-4} M$ head-on collision of a $q = 100$ binary. Snapshots are provided at the indicated times. For clarity the positions of the smaller puncture at each time have been shifted to their initial positions.*

or that no coarse puncture tracking routine can be developed, but crucially the concept itself is sound. Certainly it is worthy of further development, and with tracking high on the list of priorities we explore in the following section one proposed approach to this.

7.3. *Tracking the Puncture*

As a result of the extremely coarse grids we aim to use in the evolution of the puncture we are no longer able to accurately track its location by simply reading off the velocity from the shift vector. A new technique is required and in this section we present negative results using one potential candidate, together with a number of suggestions for improvements that could be made. We find that while the static Schwarzschild puncture can be trivially located to machine precision when it is placed at the origin, the boosted puncture deviates from its expected trajectory before crashing.

At its heart is a similar philosophy to that we have been using to evolve the puncture: using what is already known about the spacetime to inform the calculation, providing an estimate that can then be refined. With these extremely coarse resolutions that this approach relies upon the puncture is no longer covered by a mesh of grid points but sits inside a cell of side d described by the eight nearest grid points, as illustrated in FIG 7.10. We only have access to information about the spacetime at the corners of this cell surrounding the puncture, and from this alone we must determine precisely where the puncture lies. Using our analytical knowledge of the surrounding local geometry, by reading off the value of for example ϕ at a grid point we can infer how far away the puncture is. All that is required is the function $r(\phi)$, which we determine in just the same way as we found $\wp$ by utilising the previously-generated high-resolution trumpet data. The fit we use here is given by

$$r(\phi) = m \left(\frac{0.285\phi^{0.357} + 0.288\phi^{0.355} + 0.283\phi^{0.352} + 0.293\phi^{0.352} + 0.286\phi^{0.352} + 0.274\phi^{0.350}}{1.424\phi^{1.657} + 1.422\phi^{1.652} + 1.424\phi^{1.646} + 1.427\phi^{1.644} + 1.427\phi^{1.638} + 1.429\phi^{1.621}} \right),$$

7·8

 remembering to take into account mass scaling for future applications to unequal-mass binaries. But this only tells us how far away the puncture is from the grid point in absolute terms—we need to know where along the x-axis it lies. Using FIG 7.10 again for reference, this corresponds to determining the length y which can be achieved geometrically. From the forward grid point in the top-left corner we first read off the value of ϕ directly and plug it into 7·8 to find the distance $r(\phi)$. Through a simple application of Pythagoras'

theorem we use this together with the known grid spacing d to determine the length x, the distance from the grid point to the point at which the yz-face of the cell that it is on is threaded by the x-axis. Given BAM's staggered grid setup this will be, at every time step, in the very centre of the face, hence x can be calculated simply from d. With this we able to find, again through Pythagoras' theorem, the distance y between the puncture and the face of the cell. The location of the puncture is given then by

$$
\begin{aligned}
x_1 &= \text{location of grid point} + y \,, \\
&= \text{location of grid point} + \sqrt{r(\phi)^2 - 2d^2} \,.
\end{aligned}
$$

7.9

This calculation is entirely geometric and so should only depend on being correctly implemented within BAM, with the exception of the estimate of $r(\phi)$. The estimated location of the puncture is heavily dependent on the fit used here and therefore needs to be as accurate as possible. This becomes more difficult with increased distance from the puncture as the conformal factor's gradient approaches zero, meaning that it becomes increasingly difficult to distinguish one location from another; at large distances, a small variation in the value of ϕ can correspond to a large change in the resulting r. Ideally then we would make use of whichever of our variables is characterised by the greatest length scale to reduce this uncertainty as far as possible. To account for any bias that the fit to $r(\phi)$ may introduce, we perform the above calculation at each of the eight grid points surrounding the puncture and take an average of their location estimates.

In FIG 7.11 are shown the results of this tracking algorithm for both a finely-resolved simulation run with old BAM and a coarsely-resolved simulation run with new BAM, along with results using the current puncture tracker for comparison. In the case of a static Schwarzschild puncture located at the origin the new tracker performs well, identifying the location accurately to machine precision. This is not surprising given the spherical symmetry of the system and the spherically-symmetric distribution of points at which the location is calculated, as taking the average will by definition return a location of zero. More revealing is the same run with the puncture offset from the origin, here located at $x = 30M$. This no longer enjoys the same symmetry and indeed we see, first of all, the current tracking algorithm failing in coarse resolutions to determine the puncture location to the same precision as previously. But the new puncture tracker suffers from more severe issues, crashing immediately both for the old and new versions of BAM. The
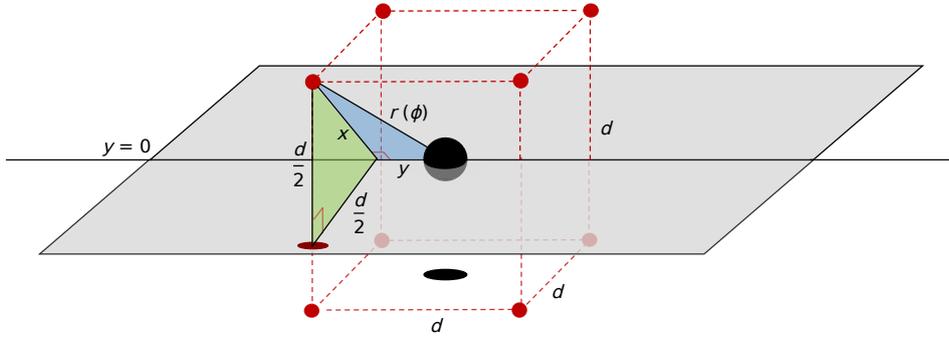
FIGURE 7.10 *Schematic of the proposed puncture-tracking mechanism for a simulation with finest grid spacing d. The puncture sits at the origin, with the x-axis seen threading the grey xy-plane in black. The red dots represent the nearest grid points to the puncture, with the projection of the puncture's location on to the base of this cell given by the black oval. The red oval is similarly the projection of the grid point on the xy-plane. The method outlined in this section involves first computing the length of side x of the green triangle followed by side y of the blue triangle using Pythagoras' theorem.*

boosted Schwarzschild case demonstrates more clearly that the current tracker is unable to accurately locate the puncture in a coarse grid, but so too is the new tracker. In the coarse grid using new BAM the simulation crashes after $45M$, and even for a finely-resolved old BAM run the trajectory soon departs from the correct values.

At the time step where the tracker fails for the offset static Schwarzschild runs we find that the reading of the conformal factor at the nearest grid points returns an undefined value. This is in a way encouraging; the reading is taken before any other calculations are performed and should be relatively trivial, so it may well be that the algorithm introduced here functions correctly if given accurate values of the conformal factor. The priority then is to ensure these values are taken accurately. Potential problems anticipated beyond this include the use of the conformal factor to estimate the distance—a more robust approach may be to use the invariant quantity $A_{ij}A^{ij}$ instead. Additionally it could well be that whichever quantity is used decays to flat space too close to the puncture. Given the coarse grid spacings that are at the heart of this approach this would prevent the accurate distinction of one location from another in the fit to $r(\phi)$. The Brill-Lindquist conformal factor for example, $\phi_{BL} = \log\left(1 + \frac{m}{2r}\right)$, only varies by $\mathcal{O}\left(10^{-4}\right)$ between distances of $20m$ and $20.1m$. Is this too small a difference to extract a precise location? Can a fit to $r(\phi)$
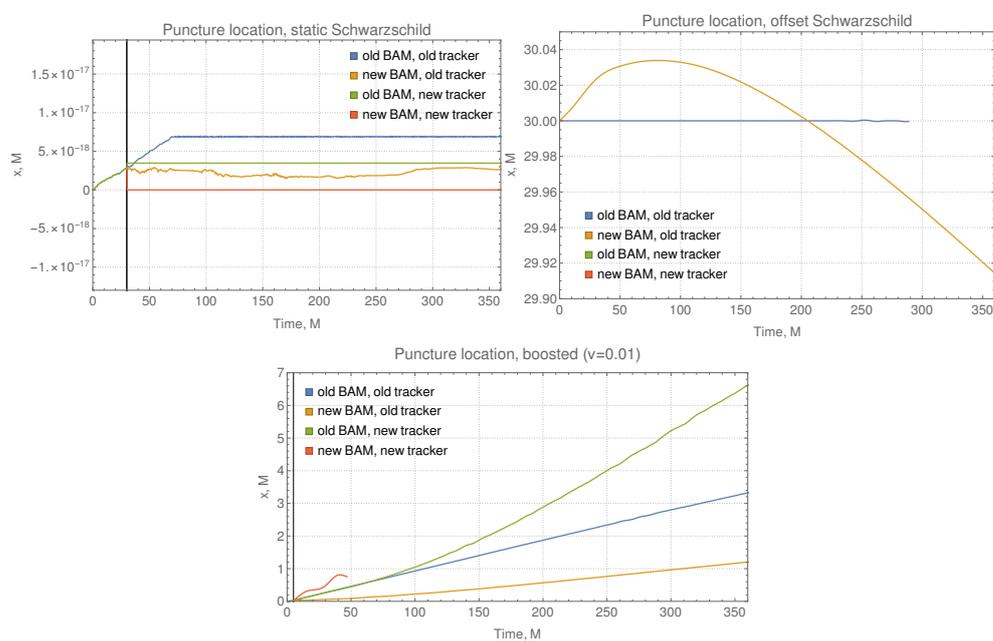
FIGURE 7.11   *Puncture location as computed using the current tracker with old BAM at $d = 4.7 \times 10^{-2}M$ (blue), the current tracker with new BAM at $d = 12M$ (yellow), the new tracker with old BAM at $d = 4.7 \times 10^{-2}M$ (green), and the new tracker with new BAM at $d = 12M$ (red). This is shown for the static Schwarzschild case in the top two panels, with the puncture located at the origin on the left and at $x = 30M$ on the right, and the boosted ($v = 0.01$) Schwarzschild case in the bottom panel.*

be constructed that is of sufficient accuracy? The impact of this will need to be analysed in the wider context of the overall error budget.

Once this procedure has been found to accurately track the offset puncture there is no reason in principle it shouldn't also work for the boosted case as long as the nearest grid points can be identified consistently when their distance from the puncture is constantly changing, and particularly as the puncture moves from one cell to another, assuming the shape of the conformal factor is not altered excessively by the boost. Turning back to FIG 7.1 this wouldn't seem to be a risk, but is worth bearing in mind nonetheless. There will however be further complications upon the introduction of a second black hole in the domain. Its effect on the conformal factor measured at each grid point must now be taken into account in the model of $r(\phi)$, not just to prevent inaccuracies in the location estimates but more importantly because the gradient leads to a negative radicand in 7.9, causing the simulation to crash.

Clearly there is at present a need for further development. For the moment though we will treat this as an issue that will at some future point be solved, and proceed with further testing of the modifications using the proxy tracker described above. This at least allows us to explore whether or not the fundamental approach to these high mass ratio simulations is viable given the correct puncture positions.

7.4. *Next Steps*

Having successfully demonstrated the stable evolution of these proof-of-principle cases in this chapter, leaving aside whether they are recovering the accurate solution or simply the phenomenology, it is clear that the approach outlined here is worth pursuing. The results presented here are of course only the preliminary stages of an extensive, and hopefully fruitful, programme of work with the end goal of fully inspiralling black hole binaries of arbitrarily high mass ratio, and as such there are a great many obstacles to overcome, both known and unknown at this time.

First, a refinement of the current implementation of these modifications. The static Schwarzschild case ran for $30,000M$; can this be extended, and if so, to what degree? Our $v = 0.01$ boosted Schwarzschild simulation started to collapse after $3000M$. Does this collapse occur sooner for higher velocities, and is there some velocity beyond which the current analytic pieces no longer accurately approximate shape of the BSSN evolution

variables? In principle there is no reason why, with sufficiently accurate analytic pieces, this method couldn't be used to stably evolve traditionally fine resolutions as well as the much coarse ones explored in this chapter, yet this hasn't been achieved so far. There are a number of outstanding issues and open questions to address here.

Perhaps the most pressing major issue yet to solve is the tracking of the motion of the puncture, as has been discussed briefly above. One currently unsucessful possibility has been outlined in §7.3 that with further work could be refined to the point that it accurately captures the location at each time step. Should the development of a coarse tracking algorithm prove to be intractable another option would be to place a grid point on the puncture itself after having regularised the evolution equations, or to rewrite the evolution equations locally (that is, in the puncture-approaching limit) to explicitly derive an evolution equation for the puncture dynamics. After all, this must reduce to Newtonian gravity at large binary separations and to PN corrections as this separation is reduced, so it can be imagined this may return the fully general relativistic equation of motion for the smaller puncture. These efforts will first depend on investigating how crucial the 'missing' features seen in FIG 7.9 are for the accurate evolution of the space-time: all that has been demonstrated here is the stable evolution of some initial data that results in the correct phenomenology. We stress that this is by no means the recovery of an accurate or indeed valid solution to the Einstein field equations, and that without the manual provision of puncture trajectories by the user even this phenomenology would not be recovered, but it is certainly an encouraging initial result that merits further investigation.

A number of relatively straightforward practical issues exist within the current implementation of this approach within the BAM code that will need to be addressed in future developments. For example it is currently incompatible with Berger-Oliger refinement of the time steps and is therefore limited to uniform. In all of the tests presented here there is no problem with the use of uniform time stepping given their relative simplicity. Berger-Oliger is a much more efficient choice when evolving simulations with a great many refinement levels, but the static and boosted Schwarzschild cases here can be run with just a single refinement level, and even the $q = 100$ case needs no more than ten levels. While this would certainly benefit from Berger-Oliger time steps uniform has been sufficient for initial tests, though as we move to higher mass ratios it will become a necessity.

More generally, the most significant issue posed by our approach, now that it has been shown in principle to work, is how heavily it relies upon knowledge of the local solution at the smaller black hole. Ideally this local solution would represent a boosted Kerr black hole; the example in FIG 6.1 used simply initial data for two momentarily stationary wormholes. The corresponding data for two trumpets in the 1+log gauge is not known in the far more general boosted spinning case, and without further developments towards this it will most likely prove impossible to isolate the background in as clean a fashion as in FIG 6.1.

There is also the nonlinearity of the Einstein equations to deal with. The Brill-Lindquist solution, in which the contributions from individual black holes can be added and subtracted freely, is a somewhat artificial construction seeing as it consists of momentarily stationary black holes with no gravitational radiation. In general we do not enjoy this freedom. But of course the conformal factor and extrinsic curvature are singular, and so the contribution from the larger black hole vanishes at the smaller; it may be expected then that given sufficiently accurate knowledge of the local black hole solution and the interaction terms, and with carefully considered gauge choices, it should be possible to construct some sufficiently detailed analytic piece that once more leaves the numerical piece characterised by a length scale determined by the larger black hole.

Finally, it will be necessary to construct initial data for a binary in trumpet 1+log form. For the initial test cases explored in this study it has been possible to use standard Bowen-York wormhole data as it relaxes to the trumpet topology: we can begin the evolution of the binary with standard methods before switching on the modifications described in this chapter and shedding refinement levels once the black hole has safely settled down to trumpet form. Clearly this isn't ideal, as we are still faced with the impractically high resolutions that restrict any potential speed-up during this initial period, which will be increasingly problematic at higher and higher mass ratios (and consequently finer and finer resolutions). If nothing else the size of the time steps is determined by the spacing of the finest refinement level at the start of the simulation, so by beginning at an extremely fine grid spacing we end up with unnecessarily fine time steps even after shedding.

For current-generation gravitational wave observatories, with numerical simulations of high mass ratio binaries it will be possible to sample previously inaccessible regions of the parameter space, and from this construct the template banks of waveform models that

are key to the search and parameter estimation in the LIGO-Virgo-KAGRA collaboration. While such binaries may well be rare at present, with the dramatic increase in events expected with third-generation detectors they could well become relatively commonplace and so work needs to be done now in preparation.

Though this work has focused on the applications to high mass ratios there is in principle no reason the methods described here cannot be applied to the comparable-mass regime. The background metric will no longer take on such a large length scale compared to the foreground black hole, but with sufficiently detailed construction of the analytic pieces it should be possible to capture those fine features. The door is then opened to rapid numerical simulations: not only would this allow far greater coverage of the parameter space (and so far more accurate waveform models), but it can be used to produce long-term simulations that capture much more of the inspiral, starting at much lower frequencies than is currently practical. Early inspiral is currently the domain of PN and EOB methods, so it would be possible to quantify the accuracy of these approximations and even use NR to tune their higher order terms. Inaccuracies in these approximations are expected to be the dominant source of systematic error in third-generation measurements as the majority of the signal power comes from the extremely long-duration in-band inspiral. Of course simply using higher order PN terms would help here, but with each successive order the calculation of these terms requires more and more work—which could well become a limiting factor in the output of future detectors.

Current efforts to produce EMRI waveforms focus on perturbative self-force calculations. The frontier in this field is the extension of these calculations to second order for generic orbits with sufficient accuracy for LISA observations, though the prospects of achieving this in the near future remain unclear [132]. The new approach introduced in this thesis opens an alternative route, one that will exploit analytic results for single black holes, binaries, and tidal terms to isolate those aspects of the numerical calculations that represent the puncture dynamics, the nonlinear interaction terms, and the gravitational wave signal. If it were computationally competitive with self-force codes this approach would enjoy the significant advantage of solving the full nonlinear Einstein equations without approximations. One illustrative example of the issues plaguing such approximations are the resonances mentioned in §3.5. The discontinuity this produces in the phase cannot be captured accurately by self-force efforts yet may be a generic feature of EMRIs, so the fully general-relativistic evolution could be required either as an entirely separate

approach to the problem or to inform self-force calculations. If EMRI waveforms can be achieved, a programme of work analogous to that used for LIGO-Virgo-KAGRA can be undertaken to sample the EMRI parameter space and develop waveform template banks in preparation for space-based detectors such as LISA, Taiji, and TianQin.

Concluding Remarks

8

The field of gravitational wave astronomy is now firmly established. In the six years since their first detection we have relied heavily on the advances made over the preceding century in anticipation of this moment: numerical relativity, and computing power along with it, had 60 years to mature to the point of simulating a full binary inspiral; post-Newtonian techniques have been continually refined ever since de Sitter, Lorentz, and Droste in the mid-1910s; waveform models and data analysis tools were all prepared in advance before a single gravitational wave had been detected. With its arrival we have been able to reap the fruits of all that labour, and finally the full force of those hundred years of preparation can be brought to bear on real astrophysical signals in order to observe the universe.

But we're rapidly approaching the limits of these 20$^{\text{th}}$ century developments. The 2020s will see the completion of the fourth and fifth LVK observing runs and the introduction of a sixth detector in the form of LIGO India, while the 2030s will bring both the third generation of ground-based detectors as well as the first space-based detectors. Our practical capabilities to handle the events they promise are severely challenged—the sheer number of anticipated signals, particularly as they are likely to overlap with one another, pose significant problems for current data analysis tools [133], and as we have focused on in this thesis current modelling efforts are insufficient given the high SNRs expected and their positions within more extreme regions of the parameter space.

The catalogue in CHAPTER 4 is one example of efforts towards future-proofing current LVK analysis workflows, contributing simulations in these under-explored parts of the parameter space to the public repository of numerical waveforms, with mass ratios

up to 8 and spins reaching 0.8. Already these have been used to improve the accuracy of existing waveform models in the precessing regime [61]. But fundamentally this accuracy is limited by that of the simulations from which these models are constructed, and in CHAPTER 5 we have found this to be somewhat lacking. From a cross-code validation study between the numerical relativity codes BAM, SpEC, and LazEv it has been shown that there exist significant discrepancies between the phases of moving-puncture and excision codes, on the order of π radians in the $(2, 2)$ mode of the strain over ten orbits. We have seen that this discrepancy will manifest itself in the analysis of astrophysical signals even with current-generation detectors for sufficiently strong signals—SNR≈20 even in optimistic cases where only the $(2, 2)$ mode is considered, comparable to the strength of the loudest black hole binary events already detected. But perhaps more concerning is the implication for third-generation detectors. It is expected that events in Einstein Telescope will regularly exceed SNRs of 100, and with thousands of events per year these phase uncertainties pose a significant risk to parameter estimation efforts as things currently stand. Improvements in systematics and conventions are required, with a greater need for communication between the research groups within the LVK that are generating numerical waveforms. The particular sources of error in the work presented here come down to a lack of precision in the specification of orbital frequency, limits on orbital eccentricity that are far too generous, and most significantly potentially inconsistent definitions of mass or choices of gauge.

We have also seen how, even more worryingly than these finer points of waveform accuracy, we currently face the complete inability to accurately model the waveform generated by one of the most anticipated candidate sources in space-based detectors: extreme mass ratio black hole binaries. Rather than some flaw in our theory or understanding of the physics this is, frustratingly, nothing more than a matter of computing power and in CHAPTER 6 we introduced a novel approach to tackling this problem. Preliminary tests on a single stationary Schwarzschild black hole offer the first hints of its viability as a solution, capable in its current form of accurately evolving black hole spacetimes in resolutions almost three orders of magnitude lower than currently required by BAM, from a grid spacing $d = 4.7 \times 10^{-2} M$ to $d = 40 M$. This represents a speed-up of 60,000%, but more importantly if it can be extended to the full binary inspiral case would allow the simulation of extreme mass ratios on a timescale comparable to current equal-mass simulations. To this end we have presented in CHAPTER 7 the extension of these pre-

liminary results to the boosted case and the first attempts at the head-on collision of a $q = 100$ binary, with resolutions of $0.2M = 20.2m_1$ demonstrated to be stable, but this has relied upon manual prescription of the puncture trajectories and neglects significant features of the full solution. Precisely how this impacts the accuracy of the resulting simulations, or indeed the viability of the approach as a whole, remains to be seen. This represents just the beginning of what will hopefully prove to be a fruitful programme of work culminating in extreme mass ratio simulations of fully inspiralling Kerr black holes in practical timescales. This will require a number of further developments though, most pressingly a way to track the location of the black hole. Aside from this, the first major hurdle to overcome would be the fashioning of far more accurate analytic pieces through the construction of 1+log binary Kerr trumpet initial data. Ultimately the goal would be to mirror the workflow established for ground-based detectors at equal-mass: build up a repository of simulations at strategic points in the parameter space, use them to inform the construction of waveform models, and apply these models to search and parameter estimation efforts by space-based detectors. It is hoped that the motion of the smaller black hole through the near-horizon geometry of the larger over the course of potentially thousands of orbits in the detector bandwidth [9] will permit precision tests of general relativity.

Bibliography

———— ❖ ————

[1] A. Einstein, "Näherungsweise Integration der Feldgleichungen der Gravitation," *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)*, pp. 688–696, Jan. 1916.

[2] D. Kennefick, *Traveling at the Speed of Thought: Einstein and the Quest for Gravitational Waves*. Princeton University Press, 2007.

[3] B. P. Abbott *et al.*, "GW150914: First results from the search for binary black hole coalescence with Advanced LIGO," *Phys. Rev. D*, vol. 93, no. 12, p. 122003, 2016.

[4] R. H. Dicke, P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson, "Cosmic Black-Body Radiation.," *ApJ.*, vol. 142, pp. 414–419, July 1965.

[5] Q. R. Ahmad *et al.*, "Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^{8}B solar neutrinos at the Sudbury Neutrino Observatory," *Phys. Rev. Lett.*, vol. 87, p. 071301, 2001.

[6] A. Królak and M. Patil, "The First Detection of Gravitational Waves," *Universe*, vol. 3, no. 3, p. 59, 2017.

[7] B. P. Abbott *et al.*, "Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO, Advanced Virgo and KAGRA," *Living Rev. Rel.*, vol. 21, no. 1, p. 3, 2018.

[8] M. Maggiore *et al.*, "Science Case for the Einstein Telescope," *JCAP*, vol. 03, p. 050, 2020.

[9] P. Amaro-Seoane, C. F. Sopuerta, and M. D. Freitag, "The role of the supermassive black hole spin in the estimation of the EMRI event rate," *Mon. Not. Roy. Astron. Soc.*, vol. 429, no. 4, pp. 3155–3165, 2013.

[10] M. Alcubierre, *Introduction to 3+1 numerical relativity*. International series of monographs on physics, Oxford: Oxford Univ. Press, 2008.

[11] S. G. Hahn and R. W. Lindquist, "The two-body problem in geometrodynamics," *Annals of Physics*, vol. 29, no. 2, pp. 304–331, 1964.

[12] R. Jantzen and P. Carini, *Understanding spacetime splittings and their relationships*, pp. 185–241. Naples: Bibliopolis, 1991.

[13] T. Baumgarte and S. Shapiro, *Numerical Relativity: Solving Einstein's Equations on the Computer*. Cambridge University Press, 2010.

[14] T. Nakamura, K. Oohara, and Y. Kojima, "General Relativistic Collapse to Black Holes and Gravitational Waves from Black Holes," *Progress of Theoretical Physics Supplement*, vol. 90, pp. 1–218, 01 1987.

[15] M. Shibata and T. Nakamura, "Evolution of three-dimensional gravitational waves: Harmonic slicing case," *Phys. Rev. D*, vol. 52, pp. 5428–5444, 1995.

[16] T. W. Baumgarte and S. L. Shapiro, "On the numerical integration of Einstein's field equations," *Phys. Rev. D*, vol. 59, p. 024007, 1998.

[17] J. W. York, Jr., "Conformal 'thin sandwich' data for the initial-value problem," *Phys. Rev. Lett.*, vol. 82, pp. 1350–1353, 1999.

[18] H. P. Pfeiffer and J. W. York, Jr., "Extrinsic curvature and the Einstein constraints," *Phys. Rev. D*, vol. 67, p. 044022, 2003.

[19] J. M. Martí and E. Müller, "The analytical solution of the riemann problem in relativistic hydrodynamics," *Journal of Fluid Mechanics*, vol. 258, p. 317–333, 1994.

[20] A. Lichnerowicz, "Sur les équations relativistes de la gravitation," *Bulletin de la Société Mathématique de France*, vol. 80, pp. 237–251, 1952.

[21] N. Ó Murchadha and J. W. York, "Initial-value problem of general relativity. 1. General formulation and physical interpretation," *Phys. Rev. D*, vol. 10, pp. 428–436, 1974.

[22] D. R. Brill and R. W. Lindquist, "Interaction energy in geometrostatics," *Phys. Rev.*, vol. 131, pp. 471–476, Jul 1963.

[23] R. W. Lindquist, "Initial-value problem on einstein-rosen manifolds," *Journal of Mathematical Physics*, vol. 4, no. 7, pp. 938–950, 1963.

[24] N. Ó Murchadha, *Boundary Conditions for the Momentum Contraint*, pp. 83–93. Cambridge University Press, 1992.

[25] S. Brandt and B. Brügmann, "A Simple construction of initial data for multiple black holes," *Phys. Rev. Lett.*, vol. 78, pp. 3606–3609, 1997.

[26] M. Alcubierre, B. Brügmann, P. Diener, M. Koppitz, D. Pollney, E. Seidel, and R. Takahashi, "Gauge conditions for long term numerical black hole evolutions without excision," *Phys. Rev. D*, vol. 67, p. 084023, 2003.

[27] M. Alcubierre, G. Allen, B. Brügmann, T. Dramlitsch, J. A. Font, P. Papadopoulos, E. Seidel, N. Stergioulas, W.-M. Suen, and R. Takahashi, "Towards a stable numerical evolution of strongly gravitating systems in general relativity: The Conformal treatments," *Phys. Rev. D*, vol. 62, p. 044034, 2000.

[28] M. Alcubierre, B. Brügmann, D. Pollney, E. Seidel, and R. Takahashi, "Black hole excision for dynamic black holes," *Phys. Rev. D*, vol. 64, p. 061501, 2001.

[29] B. Brügmann, W. Tichy, and N. Jansen, "Numerical simulation of orbiting black holes," *Phys. Rev. Lett.*, vol. 92, p. 211101, 2004.

[30] B. Brügmann, J. A. Gonzalez, M. Hannam, S. Husa, U. Sperhake, and W. Tichy, "Calibration of Moving Puncture Simulations," *Phys. Rev. D*, vol. 77, p. 024027, 2008.

[31] J. Healy, C. O. Lousto, Y. Zlochower, and M. Campanelli, "The RIT binary black hole simulations catalog," *Class. Quant. Grav.*, vol. 34, no. 22, p. 224001, 2017.

[32] K. Jani, J. Healy, J. A. Clark, L. London, P. Laguna, and D. Shoemaker, "Georgia Tech Catalog of Gravitational Waveforms," *Class. Quant. Grav.*, vol. 33, no. 20, p. 204001, 2016.

[33] C. Bona, J. Masso, E. Seidel, and J. Stela, "A New formalism for numerical relativity," *Phys. Rev. Lett.*, vol. 75, pp. 600–603, 1995.

[34] M. Hannam, S. Husa, F. Ohme, B. Brugmann, and N. Ó Murchadha, "Wormholes and trumpets: The Schwarzschild spacetime for the moving-puncture generation," *Phys. Rev. D*, vol. 78, p. 064020, 2008.

[35] A. Garat and R. H. Price, "Nonexistence of conformally flat slices of the Kerr space-time," *Phys. Rev. D*, vol. 61, p. 124011, 2000.

[36] D. C. Robinson, "Uniqueness of the Kerr black hole," *Phys. Rev. Lett.*, vol. 34, pp. 905–906, 1975.

[37] L. B. Szabados, "Quasi-Local Energy-Momentum and Angular Momentum in General Relativity," *Living Rev. Rel.*, vol. 12, p. 4, 2009.

[38] D. R. Brill, "Aspects of analyticity," *Lectures held at Kazan Summer School*, 7 1995.

[39] E. Curiel, "The many definitions of a black hole," *Nature Astron.*, vol. 3, no. 1, pp. 27–34, 2019.

[40] A. Allahyari, J. T. Firouzjaee, and R. Mansouri, "Nonlinear phenomena in general relativity," *Eur. Phys. J. Plus*, vol. 133, no. 4, p. 147, 2018.

[41] J. Thornburg, "Event and apparent horizon finders for 3+1 numerical relativity," *Living Rev. Rel.*, vol. 10, p. 3, 2007.

[42] O. Dreyer, B. Krishnan, D. Shoemaker, and E. Schnetter, "Introduction to isolated horizons in numerical relativity," *Phys. Rev.*, vol. D67, p. 024018, 2003.

[43] C. M. Werneth, M. Dhar, K. Maung Maung, C. Sirola, and J. W. Norbury, "Numerical Gram-Schmidt orthonormalization," *European Journal of Physics*, vol. 31, pp. 693–700, May 2010.

[44] B. Brügmann, "Binary black hole mergers in 3-d numerical relativity," *Int. J. Mod. Phys. D*, vol. 8, p. 85, 1999.

[45] M. Devilliers, B. Vincent, and I. Mnassri, "A new adaptive mesh refinement to model water flow around fishing nets," *Ocean Engineering*, vol. 113, pp. 34–43, 2016.

[46] G. L. Bryan, "Fluids in the universe: Adaptive mesh refinement in cosmology," *Computing in Science and Engg.*, vol. 1, p. 46–53, Mar. 1999.

[47] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, "Accurate evolutions of orbiting black-hole binaries without excision," *Phys. Rev. Lett.*, vol. 96, p. 111101, 2006.

[48] C. Gundlach, "Pseudospectral apparent horizon finders: An Efficient new algorithm," *Phys. Rev. D*, vol. 57, pp. 863–875, 1998.

[49] A. R. de Sitter, W., "On Einstein's Theory of Gravitation and its Astronomical Consequences. Second Paper.," *Monthly Notices of the Royal Astronomical Society*, vol. 77, pp. 155–184, 12 1916.

[50] H. A. Lorentz and J. Droste, *The Motion of a System of Bodies under the Influence of their Mutual Attraction, According to Einstein's Theory*, pp. 330–355. Dordrecht: Springer Netherlands, 1937.

[51] E. Hamilton and M. Hannam, "Inferring black-hole orbital dynamics from numerical-relativity gravitational waveforms," *Phys. Rev. D*, vol. 98, no. 8, p. 084018, 2018.

[52] L. Blanchet, "Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries," *Living Rev. Rel.*, vol. 17, p. 2, 2014.

[53] S. Chandrasekhar, *The Mathematical Theory of Black Holes*. International series of monographs on physics, Oxford University Press, 1992.

[54] M. Sasaki and H. Tagoshi, "Analytic black hole perturbation approach to gravitational radiation," *Living Rev. Rel.*, vol. 6, p. 6, 2003.

[55] T. Damour and A. Nagar, "The Effective One Body description of the Two-Body problem," *Fundam. Theor. Phys.*, vol. 162, pp. 211–252, 2011.

[56] M. Hannam, "Modelling gravitational waves from precessing black-hole binaries: Progress, challenges and prospects," *Gen. Rel. Grav.*, vol. 46, p. 1767, 2014.

[57] P. Schmidt, M. Hannam, S. Husa, and P. Ajith, "Tracking the precession of compact binaries from their gravitational-wave signal," *Phys. Rev. D*, vol. 84, p. 024046, 2011.

[58] R. O'Shaughnessy, B. Vaishnav, J. Healy, and D. Shoemaker, "Intrinsic selection biases of ground-based gravitational wave searches for high-mass BH-BH mergers," *Phys. Rev. D*, vol. 82, p. 104006, 2010.

[59] M. Boyle, R. Owen, and H. P. Pfeiffer, "A geometric approach to the precession of compact binaries," *Phys. Rev. D*, vol. 84, p. 124011, 2011.

[60] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, "Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms," *Phys. Rev. Lett.*, vol. 113, no. 15, p. 151101, 2014.

[61] Hamilton, Eleanor and London, Lionel and Thompson, Jonathan and Fauchon-Jones, Edward and Hannam, Mark and Hoy, Charlie and Kalaghatgi, Chinmay and Khan, Sebastian and Pannarale, Francesco and Vañó-Viñuales, Alex and Yee-les, Dave, "The final twist: A model of binary-black-hole gravitational-waveform precession geometry through merger and ringdown." In preparation, 2021.

[62] E. Berti, V. Cardoso, J. A. Gonzalez, U. Sperhake, M. Hannam, S. Husa, and B. Brügmann, "Inspiral, merger and ringdown of unequal mass black hole binaries: A Multipolar analysis," *Phys. Rev. D*, vol. 76, p. 064034, 2007.

[63] L. London, S. Khan, E. Fauchon-Jones, C. García, M. Hannam, S. Husa, X. Jiménez-Forteza, C. Kalaghatgi, F. Ohme, and F. Pannarale, "First higher-multipole model of gravitational waves from spinning and coalescing black-hole binaries," *Phys. Rev. Lett.*, vol. 120, no. 16, p. 161102, 2018.

[64] D. Dempsey, *Wave propagation on black hole spacetimes*. PhD thesis, University of Sheffield, School of Mathematics and Statistics, Sheffield, United Kindgom, 2017.

[65] J. Gair, N. Yunes, and C. M. Bender, "Resonances in Extreme Mass-Ratio Inspirals: Asymptotic and Hyperasymptotic Analysis," *J. Math. Phys.*, vol. 53, p. 032503, 2012.

[66] L. Speri and J. R. Gair, "Assessing the impact of transient orbital resonances," *Phys. Rev. D (accepted, awaiting publication)*, 2021.

[67] C. Kalaghatgi, *Modelling and studying gravitational waves from black-hole-binary mergers*. PhD thesis, Cardiff University Gravity Exploration Institute, Cardiff, United Kingdom, 2019.

[68] M.-H. Chen, Q.-M. Shao, and J. G. Ibrahim, *Monte Carlo Methods in Bayesian Computation*. Springer Series in Statistics, Springer, New York, NY, 2000.

[69] A. Taracchini *et al.*, "Effective-one-body model for black-hole binaries with generic mass ratios and spins," *Phys. Rev. D*, vol. 89, no. 6, p. 061502, 2014.

[70] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. Jiménez Forteza, and A. Bohé, "Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal," *Phys. Rev. D*, vol. 93, no. 4, p. 044006, 2016.

[71] S. Babak, A. Taracchini, and A. Buonanno, "Validating the effective-one-body model of spinning, precessing binary black holes against numerical relativity," *Phys. Rev. D*, vol. 95, no. 2, p. 024010, 2017.

[72] A. Bohé *et al.*, "Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors," *Phys. Rev. D*, vol. 95, no. 4, p. 044028, 2017.

[73] R. Cotesta, A. Buonanno, A. Bohé, A. Taracchini, I. Hinder, and S. Ossokine, "Enriching the Symphony of Gravitational Waves from Binary Black Holes by Tuning Higher Harmonics," *Phys. Rev. D*, vol. 98, no. 8, p. 084028, 2018.

[74] S. Ossokine *et al.*, "Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation," *Phys. Rev. D*, vol. 102, no. 4, p. 044055, 2020.

[75] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, "Surrogate models for precessing binary black hole simulations with unequal masses," *Phys. Rev. Research.*, vol. 1, p. 033015, 2019.

[76] D. Williams, I. S. Heng, J. Gair, J. A. Clark, and B. Khamesra, "Precessing numerical relativity waveform surrogate model for binary black holes: A Gaussian process regression approach," *Phys. Rev. D*, vol. 101, no. 6, p. 063011, 2020.

[77] N. E. M. Rifat, S. E. Field, G. Khanna, and V. Varma, "Surrogate model for gravitational wave signals from comparable and large-mass-ratio black hole binaries," *Phys. Rev. D*, vol. 101, no. 8, p. 081502, 2020.

[78] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, L. E. Kidder, and H. P. Pfeiffer, "Surrogate model of hybridized numerical relativity binary black hole waveforms," *Phys. Rev. D*, vol. 99, no. 6, p. 064045, 2019.

[79] M. Boyle *et al.*, "The SXS Collaboration catalog of binary black hole simulations," *Class. Quant. Grav.*, vol. 36, no. 19, p. 195006, 2019.

[80] B. P. Abbott *et al.*, "GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs," *Phys. Rev. X*, vol. 9, no. 3, p. 031040, 2019.

[81] R. Abbott *et al.*, "GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object," *Astrophys. J. Lett.*, vol. 896, no. 2, p. L44, 2020.

[82] E. Fauchon-Jones, E. Hamilton, M. Hannam, C. Hoy, C. Kalaghatgi, S. Khan, L. London, A. Vañó Viñuales, J. Thompson, and D. Yeeles, "BAM catalogue of binary black hole simulations." In preparation, 2021.

[83] M. Purrer, S. Husa, and M. Hannam, "An Efficient iterative method to reduce eccentricity in numerical-relativity simulations of compact binary inspiral," *Phys. Rev. D*, vol. 85, p. 124051, 2012.

[84] S. Husa, M. Hannam, J. A. Gonzalez, U. Sperhake, and B. Brügmann, "Reducing eccentricity in black-hole binary evolutions with initial parameters from post-Newtonian inspiral," *Phys. Rev. D*, vol. 77, p. 044037, 2008.

[85] E. Hamilton, *Towards a precision description of precessing black-hole-binaries*. PhD thesis, Cardiff University Gravity Exploration Institute, Cardiff, United Kingdom, 2020.

[86] S. Khan, F. Ohme, K. Chatziioannou, and M. Hannam, "Including higher order multipoles in gravitational-wave models for precessing binary black holes," *Phys. Rev. D*, vol. 101, no. 2, p. 024056, 2020.

[87] R. Abbott *et al.*, "GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses," *Phys. Rev. D*, vol. 102, no. 4, p. 043015, 2020.

[88] J. G. Baker, M. Campanelli, F. Pretorius, and Y. Zlochower, "Comparisons of binary black hole merger waveforms," *Class. Quant. Grav.*, vol. 24, pp. S25–S31, 2007.

[89] M. Hannam *et al.*, "The Samurai Project: Verifying the consistency of black-hole-binary waveforms for gravitational-wave detection," *Phys. Rev.*, vol. D79, p. 084025, 2009.

[90] G. Lovelace *et al.*, "Modeling the source of GW150914 with targeted numerical-relativity simulations," *Class. Quant. Grav.*, vol. 33, no. 24, p. 244002, 2016.

[91] R. Abbott *et al.*, "GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run," 10 2020.

[92] L. E. Kidder and L. S. Finn, "Spectral methods for numerical relativity: The Initial data problem," *Phys. Rev. D*, vol. 62, p. 084026, 2000.

[93] P. Schmidt, F. Ohme, and M. Hannam, "Towards models of gravitational waveforms from generic binaries II: Modelling precession effects with a single effective precession parameter," *Phys. Rev. D*, vol. 91, no. 2, p. 024043, 2015.

[94] C. Devine, Z. B. Etienne, and S. T. McWilliams, "Optimizing spinning time-domain gravitational waveforms for Advanced LIGO data analysis," *Class. Quant. Grav.*, vol. 33, no. 12, p. 125025, 2016.

[95] P. Ajith *et al.*, "A Template bank for gravitational waveforms from coalescing binary black holes. I. Non-spinning binaries," *Phys. Rev. D*, vol. 77, p. 104017, 2008. [Erratum: Phys.Rev.D 79, 129901 (2009)].

[96] P. Ajith *et al.*, "Inspiral-merger-ringdown waveforms for black-hole binaries with non-precessing spins," *Phys. Rev. Lett.*, vol. 106, p. 241101, 2011.

[97] L. Santamaria *et al.*, "Matching post-Newtonian and numerical relativity waveforms: systematic errors and a new phenomenological model for non-precessing black hole binaries," *Phys. Rev. D*, vol. 82, p. 064016, 2010.

[98] M. Pürrer, "Frequency domain reduced order model of aligned-spin effective-one-body waveforms with generic mass-ratios and spins," *Phys. Rev. D*, vol. 93, no. 6, p. 064041, 2016.

[99] P. Schmidt, I. W. Harry, and H. P. Pfeiffer, "Numerical Relativity Injection Infrastructure," 3 2017.

[100] L. Trefethen and M. Trummer, "An instability phenomenon in spectral methods," *SIAM Journal on Numerical Analysis*, vol. 24, pp. 1008–1023, 1987.

[101] B. Szilagyi, L. Lindblom, and M. A. Scheel, "Simulations of Binary Black Hole Mergers Using Spectral Methods," *Phys. Rev. D*, vol. 80, p. 124010, 2009.

[102] H. P. Pfeiffer, D. A. Brown, L. E. Kidder, L. Lindblom, G. Lovelace, and M. A. Scheel, "Reducing orbital eccentricity in binary black hole simulations," *Class. Quant. Grav.*, vol. 24, pp. S59–S82, 2007.

[103] A. Buonanno, L. E. Kidder, A. H. Mroue, H. P. Pfeiffer, and A. Taracchini, "Reducing orbital eccentricity of precessing black-hole binaries," *Phys. Rev. D*, vol. 83, p. 104034, 2011.

[104] A. H. Mroue and H. P. Pfeiffer, "Precessing Binary Black Holes Simulations: Quasicircular Initial Data," 10 2012.

[105] G. B. Cook and H. P. Pfeiffer, "Excision boundary conditions for black hole initial data," *Phys. Rev. D*, vol. 70, p. 104016, 2004.

[106] M. Hannam, S. Husa, F. Ohme, D. Muller, and B. Brügmann, "Simulations of black-hole binaries with unequal masses or nonprecessing spins: Accuracy, physical properties, and comparison with post-Newtonian results," *Phys. Rev. D*, vol. 82, p. 124008, 2010.

[107] H. Nakano, J. Healy, C. O. Lousto, and Y. Zlochower, "Perturbative extraction of gravitational waveforms generated with Numerical Relativity," *Phys. Rev. D*, vol. 91, no. 10, p. 104022, 2015.

[108] P. C. Peters and J. Mathews, "Gravitational radiation from point masses in a Keplerian orbit," *Phys. Rev.*, vol. 131, pp. 435–439, 1963.

[109] K. D. Kokkotas and B. G. Schmidt, "Quasinormal modes of stars and black holes," *Living Rev. Rel.*, vol. 2, p. 2, 1999.

[110] A. Buikema *et al.*, "Sensitivity and performance of the Advanced LIGO detectors in the third observing run," *Phys. Rev. D*, vol. 102, no. 6, p. 062003, 2020.

[111] C. O. Lousto and Y. Zlochower, "Black hole binary remnant mass and spin: A new phenomenological formula," *Phys. Rev. D*, vol. 89, no. 10, p. 104052, 2014.

[112] M. Boyle, L. E. Kidder, S. Ossokine, and H. P. Pfeiffer, "Gravitational-wave modes from precessing black-hole binaries," 9 2014.

[113] A. Ghosh, N. K. Johnson-Mcdaniel, A. Ghosh, C. K. Mishra, P. Ajith, W. Del Pozzo, C. P. L. Berry, A. B. Nielsen, and L. London, "Testing general relativity using gravitational wave signals from the inspiral, merger and ringdown of binary black holes," *Class. Quant. Grav.*, vol. 35, no. 1, p. 014002, 2018.

[114] P. Jaranowski and A. Krolak, *Analysis of Gravitational-Wave Data*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology, Cambridge University Press, 2009.

[115] M. Hannam, S. Husa, B. Brügmann, and A. Gopakumar, "Comparison between numerical-relativity and post-Newtonian waveforms from spinning binaries: The Orbital hang-up case," *Phys. Rev. D*, vol. 78, p. 104007, 2008.

[116] L. Barsotti, S. Gras, M. Evans, and P. Fritschel, "Ligo-t1800044-v5: The updated advanced ligo design curve," tech. rep., LIGO Scientific Collaboration, 04 2018.

[117] E. Baird, S. Fairhurst, M. Hannam, and P. Murphy, "Degeneracy between mass and spin in black-hole-binary waveforms," *Phys. Rev.*, vol. D87, no. 2, p. 024035, 2013.

[118] L. S. Finn and D. F. Chernoff, "Observing binary inspiral in gravitational radiation: One interferometer," *Phys. Rev. D*, vol. 47, pp. 2198–2219, 1993.

[119] H.-Y. Chen and D. E. Holz, "The Loudest Gravitational Wave Events," 9 2014.

[120] M. Alcubierre, "Hyperbolic slicings of space-time: Singularity avoidance and gauge shocks," *Class. Quant. Grav.*, vol. 20, pp. 607–624, 2003.

[121] K. S. Thorne, "Gravitational waves from compact bodies," *IAU Symp.*, vol. 165, p. 153, 1996.

[122] J. R. Gair, M. Vallisneri, S. L. Larson, and J. G. Baker, "Testing General Relativity with Low-Frequency, Space-Based Gravitational-Wave Detectors," *Living Rev. Rel.*, vol. 16, p. 7, 2013.

[123] M. W. Choptuik, "Experiences with an adaptive mesh refinement algorithm in numerical relativity," *Frontiers in Numerical Relativity*, pp. 206–221, 1989.

[124] C. O. Lousto and Y. Zlochower, "Orbital Evolution of Extreme-Mass-Ratio Black-Hole Binaries with Numerical Relativity," *Phys. Rev. Lett.*, vol. 106, p. 041101, 2011.

[125] C. O. Lousto and J. Healy, "Exploring the Small Mass Ratio Binary Black Hole Merger via Zeno's Dichotomy Approach," *Phys. Rev. Lett.*, vol. 125, no. 19, p. 191102, 2020.

[126] M. Hannam, S. Husa, U. Sperhake, B. Brügmann, and J. A. Gonzalez, "Where post-Newtonian and numerical-relativity waveforms meet," *Phys. Rev. D*, vol. 77, p. 044020, 2008.

[127] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C (2nd Ed.): The Art of Scientific Computing*. USA: Cambridge University Press, 1992.

[128] M. Alcubierre *et al.*, "Toward standard testbeds for numerical relativity," *Class. Quant. Grav.*, vol. 21, no. 2, pp. 589–613, 2004.

[129] P. Figueras, M. Kunesch, and S. Tunyasuvunakool, "End Point of Black Ring Instabilities and the Weak Cosmic Censorship Conjecture," *Phys. Rev. Lett.*, vol. 116, no. 7, p. 071102, 2016.

[130] A. P. Gentle, D. E. Holz, A. Kheyfets, P. Laguna, W. A. Miller, and D. M. Shoemaker, "Constant crunch coordinates for black hole simulations," *Phys. Rev. D*, vol. 63, p. 064024, 2001.

[131] M. D. Hannam and G. B. Cook, "Conformal thin-sandwich puncture initial data for boosted black holes," *Phys. Rev. D*, vol. 71, p. 084023, 2005.

[132] L. Barack and A. Pound, "Self-force and radiation reaction in general relativity," *Rept. Prog. Phys.*, vol. 82, no. 1, p. 016904, 2019.

[133] A. Samajdar, J. Janquart, C. Van Den Broeck, and T. Dietrich, "Biases in parameter estimation from overlapping gravitational-wave signals in the third generation detector era," 2 2021.