The Pennsylvania State University

The Graduate School

The Eberly College of Science

THE NEXT GENERATION OF BINARY BLACK HOLE HEAD-ON COLLISIONS, AND THEIR AFTERMATH

A Thesis in

Physics

by

Bernard J. Kelly

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

August 2004

The thesis of Bernard J. Kelly was reviewed and approved* by the following:

Pablo Laguna
Professor of Physics
Professor of Astronomy & Astrophysics
Thesis Adviser
Chair of Committee

Bernd Brügmann
Associate Professor of Physics

Lee Samuel Finn
Professor of Physics
Professor of Astronomy & Astrophysics

Steinn Sigurdsson
Assistant Professor of Astronomy & Astrophysics

Jayanth Banavar
Professor of Physics
Head of the Department of Physics

*Signatures are on file in the Graduate School.

Abstract

I present new results in the full 3D evolution of black hole binary systems, for the head-on collision case. In contrast to most prior work on this problem, the initial data used is not restricted to conformally flat data sets, and the curvature singularities are not locked to the same static coordinate location on the computational domain.

With my collaborators, I have succeeded in producing long-term stable evolutions using a combination of modern techniques: an adjusted BSSN evolution system with densitized lapse function; dynamic excision of singularities, with automatic re-centering; analytic pre-merger slicing conditions pre-merger; and live "driver" slicing conditions post-merger.

The resulting evolutions are the first stable mergers of Kerr-Schild-like black holes using singularity excision. I present the results of this work for two regimes: holes with small initial separation, which represent highly distorted single black holes, and holes with large initial separation, which must undergo in-fall and merger. This work is intended to complement earlier axisymmetric simulations, and recent evolutions of binaries using very different techniques.

To accompany these full 3D evolutions, I also present published work on the use of derived wave forms to test the self-consistency of general relativity. The test described will be applicable to current and future interferometric gravitational-wave observatories, including LIGO and LISA.

Table of Contents

List of Figures

List of Tables

Acknowledgments

Many people have contributed to the research within this thesis, and to my well-being outside it; to them all I am extremely grateful. In no particular order, I would like to mention the following friends and colleagues:

Pablo Laguna – my academic adviser. He gave great encouragement in steering me toward interesting problems, and showed enormous enthusiasm in our discussions of the possible. It was from Pablo that I learned that it's permissible – and fun – to disagree with faculty occasionally.

Erik Schnetter, who constructed the first Maya code for us, and who has since worked tirelessly – on his own, and with the Cactus developers – to make useful tools for us and the entire community.

Uli Sperhake, who took the mess we'd made of Erik's nice code, and painstakingly worked out how we'd buggered it up. The Maya project might well not be alive if not for his weeping and gnashing of teeth.

Ken Smith, master of singularity excision, data visualization, a good friend who doesn't mind repetition, and a good friend who doesn't mind repetition.

The other members of the Penn State NR group, and the Centers of Gravity, past and present, for supplying an intellectually engaging environment. In particular, I would like to acknowledge Abhay Ashtekar, Bernd Brügmann, Sam Finn, Deirdre Shoemaker, Steinn Sigurdsson and Jorge Pullin, for their interest in my academic and personal well-being over the last few years.

On the more personal side, I'd like to thank my family and friends in Dublin, to whom I've failed every Christmas to explain what I'm doing in Pennsylvania. Bunsen-burner jokes never get old.

Thanks also to the applied statisticians, Amanda, Badri, Ian and Patrick.

Most of all, I'd like to thank my fiancée Christine, for her patience and tolerance over the last few months, not to say years. And if she hadn't been so keen to have something to edit, I could have dragged this out another six months or so.

BRUTUS:
No more, I pray you.
Messala, I have here received letters,
That young Octavius and Mark Antony
Come down upon us with a mighty power,
Bending their expedition toward Philippi.

MESSALA:
Myself have letters of the selfsame tenor.

BRUTUS:
With what addition?

MESSALA:
That by proscription and bills of outlawry,
Octavius, Antony, and Lepidus,
Have put to death an hundred senators.

– Julius Caesar, Act IV, Scene III

Chapter 1

Introduction

The past decade has seen the birth of the field of Gravitational-Wave Phenomenology. Several groups around the world have constructed detector arrays (LIGO, TAMA, GEO), and have embarked on a series of trial runs, taking real data, and steadily working their way toward optimal sensitivity; the first science data from the LIGO project has been analyzed [43, 44, 45, 46]. There has also been steady work done on a space-based antenna – LISA – that will be able to achieve far better sensitivity, albeit in a very different frequency band.

With this experimental progress has come a demand for wave-forms – algebraic expressions that describe the gravitational radiation the gravitational-wave physicists should be looking for. Gravitational waves (GW), while certainly extremely weak, are also extremely pervasive. The universe is most likely awash in space-time ripples from countless sources of diverse nature – super-massive black holes at the center of galaxies, pulsar systems containing neutron stars and/or low-mass black holes, massive stars undergoing supernova, and perhaps exotic matter sources we have not conceived of to date.

Thus the onus has fallen to theorists to supply the templates that can distinguish these sources from one another; so far, the results have not been impressive. Perturbative schemes, describing specific limits of source behavior, have been available for years. Post-Newtonian (PN) results tell us what we should expect from well-separated compact bodies. Close-limit (CL) theory gives us perturbative results for when a dynamical situation has proceeded to the point of mild distortion of the final state of the system. What is missing is the strong-field stage of the interaction of compact, massive, bodies. This is the regime where no perturbative scheme can apply, and we need to tackle the full theory of general relativity directly. This regime is manifestly beyond the abilities of pencil-and-paper theorists; the hope that arrived with the age of supercomputers was that numerical methods might extract the missing information by brute force.

2004 is the fortieth anniversary of one of the first serious attempts to solve Einstein's equations for a binary black-hole situation. Hahn & Lindquist [79] carried out some initial investigations in axisymmetry, and encouraged by their results, stated:

"In summary, the numerical solution of the Einstein equations presents no insurmountable difficulties."

With the benefit of hindsight, this optimism was premature. Despite great progress in

the last few years, it would be fair to say that the binary-black-hole problem of numerical relativity has not been completely solved. Early work generated impressive wave forms, but only in cases of considerable symmetry, and for very restrictive types of physical situation, which might have little connection to astrophysical reality. Over the next three decades, our computational resources have increased almost exponentially[1], giving us faster and more powerful computers with which to break the back of Einstein's equations. This brute-force approach seemed to work for a long time. And yet comparatively little progress was made in full 3D evolutions of more general situations.

It was only in the late 1990s that it became widely recognized that the real bottle-neck was not computational power, but some fundamentally awkward mathematical properties of Einstein's equations. Since this realization, huge progress has finally been made in 3D simulations that begin to approach the robustness and stability of 2D axisymmetric codes of the 1970s. We are finally approaching a stage where we can extract some meaningful results from a computational code – results that might actually correspond to something astrophysically observable.

It is in this context that I would like to present the work in this thesis. My main focus has been on re-investigating the head-on collision of two equal-mass black holes. While head-on collisions are not new by any means – this was the problem addressed by Hahn & Lindquist forty years ago – the data sets I am considering have only been evolved by one group [31]. Since those investigations, advances have led to the development of stability-enhancing evolutions systems, and more sophisticated techniques for handling the physical singularities present. By use of these techniques, I have been able to obtain long-term stable evolutions of binary systems, for a range of initial coordinate separations.

This thesis is laid out as follows: In Chapter 2, I will give a brief review of the basics of Cauchy evolution of a space-time on a computer, Penn State's Maya code and related issues.

In Chapter 3, I introduce the initial data we use to represent a binary system. I discuss the procedure used to solve the initial-value constraint equations, as well as an appropriate set of slicing conditions of this data. I also introduce some diagnostic tools used in the analysis of evolutions.

In Chapters 4 and 5, I present the results of our recent reinvestigation of the binary black hole head-on collision problem, featuring the first long-term stable evolutions using Kerr-Schild data and dynamic excision. This work is a collaboration with Pablo Laguna, Kenneth Smith, and Ulrich Sperhake.

Encouraged by these results, I would like to present, in Chapter 6, a statistical investigation concerning the eventual use of wave forms from such evolutions – a new way to classify black hole sources from observed gravitational waves, or even a statistical test of

[1] Moore's "Law" of Computational Resources predicts a doubling time of 18 months for processor speeds.

the self-consistency of General Relativity itself, and in particular, the "No-Hair Theorem". This work was done in collaboration with Olaf Dreyer, Lee Samuel Finn, David Garrison, Badri Krishnan, and Ramon Lopez-Aleman, and has been published in *Class. Quantum Grav.* [58].

1.1 Conventions

Throughout this thesis, the space-time metric is taken to have signature $(-,+,+,+)$. For almost all purposes, I use units where Newton's constant of gravitation and the speed of light in vacuum are dimensionless, with value unity: $G = c = 1$.

Cartesian spatial coordinates, when they occur, will be in the order (x, y, z). Spherical polar coordinates will be in the order (r, θ, ϕ), while cylindrical polar coordinates will be in the order (q, ϕ, z). Here, I define

$$q \equiv \sqrt{x^2 + y^2} \quad , \quad r \equiv \sqrt{q^2 + z^2} = \sqrt{x^2 + y^2 + z^2}.$$

For the most part, when dealing with tensor indices, spatial indices will be denoted by lower-case Latin letters, while four-dimensional space-time indices, when they occur, will have lower-case Greek indices. The exception to this rule will be in the Appendix Section B.2, where Latin indices will be used throughout.

As the subject matter of this thesis has to do with the evolution of spatial fields, tensor fields will be assumed to be spatial unless explicitly stated otherwise. In particular, the four-dimensional Riemann and Ricci tensors will be distinguished by a prior superscript "4".

There are several sign conventions that are not fully standardized. The Riemann and Ricci tensors (in any dimension) are given by:

$$\begin{aligned} R^a_{bcd} &\equiv \partial_c \Gamma^a_{bd} - \partial_d \Gamma^a_{bc} + \Gamma^a_{ac}\Gamma^a_{bd} - \Gamma^a_{ad}\Gamma^a_{bc}, \\ R_{ab} &\equiv R^c_{acb}. \end{aligned}$$

We use York's definition of the extrinsic curvature ([141] eq. (19),(35)):

$$\begin{aligned} K_{ij} &\equiv -\frac{1}{2}\mathcal{L}_n h_{ij} \\ &= \frac{1}{2\alpha}\left(\beta_{i|j} + \beta_{j|i} - \partial_t \gamma_{ij}\right) = \frac{1}{2\alpha}\left(\mathcal{L}_\beta h_{ij} - \mathcal{L}_t h_{ij}\right) \\ &= -\alpha\, {}^4\Gamma^0_{ij} \end{aligned}$$

This agrees with Misner et al [95]. Wald [135] uses the opposite sign (eqns (10.2.13),(E.2.30)) but is self-consistent.

Another sign convention ambiguity in the definition of the Weyl scalars Ψ_i used for diagnostics purposes (see Section 3.7 and Appendix Section B.2). The confusion arises due to some authors working with spinors, and sometimes the opposite metric signature, $(+,-,-,-)$. The overall convention is not too important: the definitions I use ensure that with an adapted tetrad, the Schwarzschild solution yields a strictly positive Ψ_2 and Kretschmann invariant I.

Chapter 2

The 3+1 Split of General Relativity and the Maya Code

2.1 The Standard ADM System

In Einstein's General Theory of Relativity (GR for short), our universe is represented by a four-dimensional differential manifold, with one time dimension, and three spatial dimensions. Gravity is the natural curvature of space-time, caused by the presence of matter. This is encapsulated mathematically by Einstein's equation:

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = 8\pi T_{\mu\nu} \tag{2.1}$$

Roughly speaking, the left-hand side of (2.1) represents the curvature of space-time, and the right-hand side represents the matter content. **G** is constructed from the metric and its first and second derivatives. Thus, whatever coordinates we choose, the equations will be second-order in time, second-order in space.

GR describes *all* of space-time simultaneously: a solution to Einstein's equations tells us about the entire history of the universe at once. However, our perception is of a world that changes through time. Formulating a problem to be solved numerically also requires that we reduce it to an evolution in time. In the 1950s, Arnowitt, Deser, and Misner (ADM) expressed GR as an initial-value Cauchy problem – an approach known as *geometrodynamics* (see their 1962 review article [12]). The original ADM system arose from their research into a canonical formulation of general relativity in 3+1 dimensions that could be quantized in a manner analogous to electrodynamics. The fundamental canonical variables are:

$$\begin{aligned} \gamma_{ij} &\equiv g_{ij}, \\ \pi^{ij} &\equiv \sqrt{|g|}\left({}^4\Gamma^0_{pq} - g_{pq}\, {}^4\Gamma^0_{rs} g^{rs} \right) g^{ip} g^{jq}. \end{aligned} \tag{2.2}$$

Note that while γ_{ij} is a three-dimensional tensor, its canonical conjugate π^{ij} is by necessity a tensor density.

The original ADM variables were appropriate in the context of canonical quantum grav-

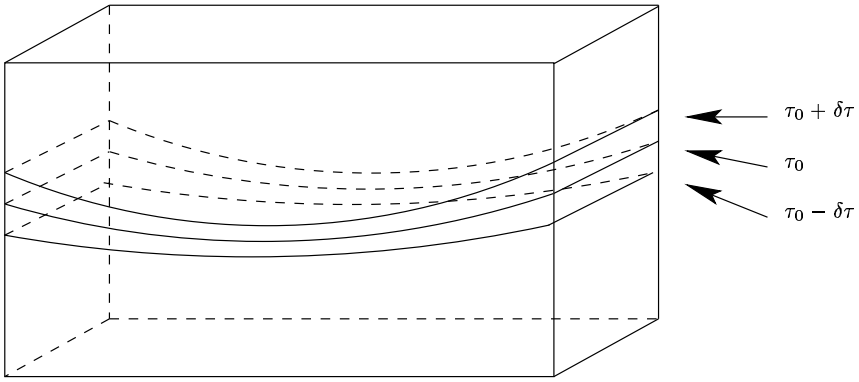

Figure 2.1: A possible foliation of space-time. Every point is contained in exactly one slice.

ity; however, for numerical purposes a simpler related system was used, referred to here as the *standard ADM* system. The definitive introduction to the standard ADM system was given by York [141].

In the standard ADM picture, space-time is sliced up (or "foliated") into a set of three-dimensional spatial hypersurfaces (Fig 2.1). These surfaces together fill the manifold without intersecting: any point in the manifold belongs to one and only one hypersurface. The intrinsic geometry of each hypersurface can be encoded in a three-dimensional metric tensor. This tensor γ – the three-metric – is our fundamental evolution variable. We want to start from an initial γ, describing some interesting physical situation (e.g., two distinct black holes), and evolve the three-metric until a later time.

Any foliation of space-time into hypersurfaces will define an unambiguous time-like unit normal to the surfaces, $\hat{\mathbf{n}}$. $\hat{\mathbf{n}}$ can be used to project all tensors onto the hypersurface, most notably, the metric itself. For instance, the matter tensor $\mathbf{T}$ can be decomposed via this projector:

$$T_{\mu\nu} = S_{\mu\nu} + j_\mu \hat{n}_\nu + j_\nu \hat{n}_\mu + \rho \hat{n}_\mu \hat{n}_\nu, \tag{2.3}$$

where ρ can be interpreted as a local energy density, j^μ as a current, and $S_{\mu\nu}$ as a local stress tensor. The left-hand side of (2.1) can be decomposed in a similar way. We can also produce a spatial covariant derivative $\mathbf{D}$, compatible with the spatial metric.

Our time direction is defined by some time-like, future-pointing vector related to $\hat{\mathbf{n}}$:

$$t^\mu = N^\mu + \beta^\mu = \alpha \hat{n}^\mu + \beta^\mu, \tag{2.4}$$

where α is a positive scalar function β^μ is a spatial vector; these are called the *lapse function* and *shift vector*, respectively. They encode some of the coordinate freedom of GR – the freedom to slice the space-time up as we see fit. Fig 2.1 illustrates their relationship to the slicing. α gives the local distance between neighboring hypersurfaces; if $\alpha \to 0$ in a region,

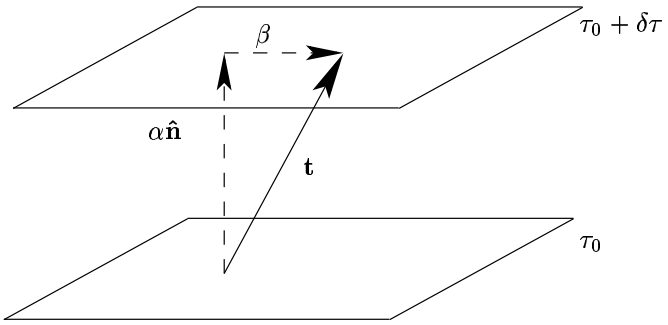

Figure 2.2: Two spatial slices connected by the lapse α and shift β^i.

the evolution "freezes" in that region. The shift β^i is so called because it encapsulates the freedom of points to move around spatially from slice to slice.

If we know the three-metric γ and the slicing fields α and β as functions of time, we can reconstruct the four-metric $\mathbf{g}$:

$$ds^2 = g_{\mu\nu}\, dx^\mu\, dx^\nu = -\alpha^2 dt^2 + \gamma_{ij}(dx^i + \beta^i dt)(dx^j + \beta^j dt) \tag{2.5}$$

Then we can write a matrix-valued equation relating the four-metric to the three-metric, lapse and shift:

$$g_{\mu\nu} = \begin{pmatrix} -\alpha^2 + \beta_p\beta^p & \beta_j \\ \beta_i & \gamma_{ij} \end{pmatrix} \Rightarrow g^{\mu\nu} = \frac{1}{\alpha^2}\begin{pmatrix} -1 & \beta^j \\ \beta^i & \gamma^{ij}\alpha^2 - \beta^i\beta^j \end{pmatrix} \tag{2.6}$$

The standard ADM system expresses Einstein's equations as a coupled set of first-order-in-time partial differential equations (PDEs). The first describes the evolution of the three-metric γ_{ij} in terms of a momentum-like quantity, the *extrinsic curvature* K_{ij}, related to the Lie derivative along $\mathbf{t}$:

$$\partial_t\, \gamma_{ij} \equiv \mathcal{L}_t\gamma_{ij} = \mathcal{L}_\beta\gamma_{ij} - 2\alpha K_{ij}. \tag{2.7}$$

The second set of equations captures the actual dynamics of the system, as it describes the evolution of K_{ij} itself:

$$\begin{aligned} \partial_t\, K_{ij} \equiv \mathcal{L}_t K_{ij} \quad &= \quad \mathcal{L}_\beta K_{ij} - D_i D_j \alpha \\ &\quad +\alpha\left[R_{ij} + KK_{ij} - 2K_{ic}K^c_{\ j} - 8\pi\left(S_{ij} - \frac{1}{2}\gamma_{ij}(S-\rho)\right)\right]. \end{aligned} \tag{2.8}$$

These evolution equations are derived from the spatial projection of (2.1). The remaining four independent parts of the Einstein equations contain no explicit time derivatives, but merely describe extra conditions that the three-metric and extrinsic curvature must obey.

These equations are constraints on the data on any spatial slice. To be specific, they are the (scalar) *Hamiltonian constraint*:

$$R + K^2 - K_{ab}K^{ab} = 16\pi\rho, \tag{2.9}$$

and the (vector) *momentum constraint*:

$$D_b \left(K^{ab} - K\gamma^{ab} \right) = 8\pi j^a. \tag{2.10}$$

This is essentially the same behavior you would expect from a problem in Hamiltonian mechanics, with the three-metric as the generalized coordinate, and the extrinsic curvature trivially related to the conjugate momentum. These four equations, (2.7), (2.8), (2.9), (2.10), are the fundamental equations of the 3+1 version of general relativity – they determine the history of the entire space-time. Given full data on an initial slice (that is, γ and $\mathbf{K}$ that solve the constraints), we may hope to find out the entire future behavior of the space. This is the fundamental mission of numerical relativity.

2.1.1 Exact Black-Hole Solutions in Numerical Relativity

Although numerical relativity is needed in generic situations where exact mathematical solutions to Einstein's equations are not known, it is certainly true that we should be able to represent known exact solutions as well. We are usually given an exact space-time via the four-metric $\mathbf{g}$. By writing down an explicit metric in specific coordinates, we are, amongst other things, imposing a slicing (that is, the lapse and shift functions). This relationship is given explicitly by (2.5).

As an example, we can express the Schwarzschild solution – the unique spherically symmetric vacuum solution to Einstein's equations – in many different ways. One way is in Schwarzschild coordinates:

$$ds^2 = -\left(1 - \frac{2M}{r}\right) dt^2 + \left(1 - \frac{2M}{r}\right)^{-1} dr^2 + r^2\, d\Omega^2. \tag{2.11}$$

However, these coordinates are not well suited to numerical work – there is a coordinate singularity at the event horizon $r = 2M$ where many of the fields become singular. As computers are not good at handling singularities, we might prefer a better-behaved form. For instance, the ingoing Eddington-Finkelstein (IEF) solution is regular at the horizon:

$$ds^2 = -\left(1 - \frac{2M}{r}\right) dt^2 + \frac{4M}{r} dt\, dr + \frac{r + 2M}{r} dr^2 + r^2 d\Omega^2. \tag{2.12}$$

This is, of course, exactly the same space-time as before, but sliced in a different way; a different lapse and shift give rise to a different time direction. In fact, the difference between

the solutions is simply the definition of a new time coordinate t:

$$t_S \equiv t_{IEF} + 2M \log\left(\frac{D}{r-2M}\right).$$

In Appendix A, I present for future convenience several exact black-hole solutions, and their decomposition into ADM fields suitable for evolution.

2.1.2 Standard ADM in Practice: Early Black Hole Evolutions

The ADM equations above were worked out in the 1960s, and when numerical simulation of space-times became a real prospect in the early 1970s, the ADM system was the obvious – and just about the only – way to consider evolving Cauchy initial data.

From the 1970s to the late 1990s, the standard ADM system was used by almost all groups, the general belief being that greater computer power would cure many of the ills seen. Only later did it appear that there were some problems with the equations as derived – mathematical problems, in the continuum – that no amount of increased computer power could solve. Ignoring finite-difference errors, the continuum Einstein field equations at any point in space can be written in the form of a quasi-linear first-order system:

$$\partial_t \vec{u} = A^i(u)\ \partial_i \vec{u} + B(u)\ \vec{u}.$$

It can be shown that the mathematical well-posedness of such a system depends on the formulation of Einstein's equations used, as this determines the eigenfunctions and eigenvalues of the system. The system can be enlarged by using a more complicated decomposition of the field equations, or by making the slicing fields α and β^i dynamic. Additionally, the eigen-structure can be altered by making algebraic gauge choices or adding constraints (which by definition are zero, and thus do not change the physical solutions of the system) to certain equations.

Many investigations have been carried out into such possible alterations. Kelly et al [83] showed that stable black-hole evolutions are possible with minor adjustments of the ADM equations in 1D using algebraic gauge conditions. Yoneda & Shinkai [138] have carried out extensive experiments with constraint addition to the ADM system.

Meanwhile, much effort has been put into developing Cauchy systems very unlike ADM, which have a well-understood characteristic structure for all fields. Among these are the Kidder-Scheel-Teukolsky [87], the Einstein-Christoffel [71] and Alekseenko-Arnold [7] systems. Not all of these have been applied to the full non-linear black hole problem, however; a proper mathematical treatment often requires very careful treatment of the characteristic fields at domain boundaries. The challenge of consistent boundary treatments is being addressed by several groups, most notably in LSU [38] and Southampton [75].

2.2 The BSSN System and Stable Single Black Hole Evolutions

By the late 1990s, it was beginning to be recognized that the standard ADM equations were ill-adapted to numerical black hole evolution schemes. In 1998 Baumgarte & Shapiro [19] reinvestigated a reformulation of the ADM equations originally suggested by Shibata & Nakamura three years earlier [116]. As this "BSSN" scheme is fast becoming the standard 3+1 Cauchy evolution system for black hole evolutions, I shall introduce it in outline, using the notation of [19].

The first step in transforming the standard ADM system is the conformal split. The physical three-metric has a non-trivial determinant, which, being strictly positive, can be expressed as an exponential:

$$|\gamma| \equiv \det(\gamma_{ij}) = e^{12\phi}. \tag{2.13}$$

We then factor this out of γ_{ij} by multiplication by the appropriate power of the determinant. The resulting "conformal metric" has unit determinant by construction:

$$\tilde{\gamma}_{ij} \equiv |\gamma|^{-\frac{1}{3}}\gamma_{ij} = e^{-4\phi}\gamma_{ij}. \tag{2.14}$$

Having made this conformal split for the three-metric, for consistency we extend it to the inverse metric, and the extrinsic curvature:

$$\begin{aligned} \tilde{\gamma}^{ij} &\equiv e^{4\phi}\gamma^{ij}, \\ \tilde{K}_{ij} &\equiv e^{-4\phi}K_{ij}. \end{aligned} \tag{2.15}$$

Note that scalar quantities, such as the trace of K_{ij}, are not affected at all by this operation.

The second step is the split of the extrinsic curvature into its trace and trace-free parts:

$$\tilde{K}_{ij} \equiv \tilde{A}_{ij} + \frac{1}{3}K\tilde{\gamma}_{ij}. \tag{2.16}$$

(In fact, this operation could have been carried out before the conformal rescaling.)

These two steps are meaningless unless something interesting happens for evolution. Both ϕ and K are elevated to the status of independent evolution variables. The final step in the creation of the BSSN system is the introduction of some variables representing spatial derivatives:

$$\tilde{\Gamma}^{i} \equiv \tilde{\gamma}^{mn}\tilde{\Gamma}^{i}_{mn}. \tag{2.17}$$

The evolution equations for these new BSSN fields can then be worked out from the standard

ADM ones:

$$\partial_t\ \tilde{\gamma}_{ij} = \mathcal{L}_\beta\ \tilde{\gamma} ij - 2\alpha\tilde{A}_{ij}, \tag{2.18}$$

$$\partial_t\ \phi = \mathcal{L}_\beta\ \phi - \frac{1}{6}\alpha K, \tag{2.19}$$

$$\begin{aligned}\partial_t\ \tilde{A}_{ij} &= \mathcal{L}_\beta\ \tilde{A}_{ij}e^{-4\phi}[-D_iD_j\alpha + \alpha R_{ij}]^{TF} \\ &\quad + \alpha(K\tilde{A}_{ij} - 2\tilde{A}_{ik}\tilde{A}^k{}_j),\end{aligned} \tag{2.20}$$

$$\begin{aligned}\partial_t\ K &= \mathcal{L}_\beta\ K - D^iD_i\alpha \\ &\quad + \alpha(\tilde{A}_{ij}\tilde{A}^{ij} + \frac{1}{3}K^2),\end{aligned} \tag{2.21}$$

$$\begin{aligned}\partial_t\ \tilde{\Gamma}^i &= \tilde{\gamma}^{jk}\partial_j\partial_k\beta^i + \frac{1}{3}\tilde{\gamma}^{ij}\partial_j\partial_k\beta^k + \beta^j\partial_j\tilde{\Gamma}^i - \tilde{\Gamma}^j\partial_j\beta^i + \frac{2}{3}\tilde{\Gamma}^i\partial_j\beta^j \\ &\quad -2\tilde{A}^{ij}\partial_j\alpha + 2\alpha(\tilde{\Gamma}^i{}_{jk}\tilde{A}^{jk} + 6\tilde{A}^{ij}\partial_j\phi - \frac{2}{3}\tilde{\gamma}^{ij}\partial_j K).\end{aligned} \tag{2.22}$$

It should be noted here that for tensor densities such as $\tilde{\gamma}_{ij}$ and $\tilde{A}_{ij}$, the Lie derivative terms are adjusted by the density weight:

$$\mathcal{L}_\beta\ \tilde{A}_{ij} \equiv \beta^k\partial_k\tilde{A}_{ij} + \tilde{A}_{ik}\partial_j\beta^k + \tilde{A}_{jk}\partial_i\beta^k - \frac{2}{3}\tilde{A}_{ij}\partial_k\beta^k.$$

Similarly, the field ϕ is not a true scalar, nor even a scalar density (though e^ϕ would be), and the $\mathcal{L}_\beta\ \phi$ in (2.19) is just defined to be that combination of terms needed to make the evolution equations for $\tilde{\gamma}_{ij}$ and ϕ consistent with the standard ADM evolution equation for γ_{ij} (2.7).

The effect of this decomposition has been to produce greatly improved evolution times in black-hole space-times (and in matter situations, also). Why exactly this should be so is still the subject of investigation. Baumgarte & Shapiro do, however, note two mathematically desirable features of the system as a justification for the transformations. Firstly, the conformal variables, $\tilde{\gamma}_{ij}$ and $\tilde{A}_{ij}$, should encapsulate all the radiation content of the space-time [140]. Secondly, the introduction of $\tilde{\Gamma}^i$ transforms the spatial part of $\partial_t\ \tilde{A}_{ij}$, yielding an explicitly elliptic operator. The result is that the BSSN system contains a coupled set of wave equations for $\tilde{\gamma}_{ij}$ and $\tilde{A}_{ij}$.

More recently, Gentle et al [73] have pointed out that the standard way of writing the evolution equation for $\tilde{\Gamma}^i$ implicitly embodies the momentum constraint, thus leading to a more constraint-preserving system.

The most crucial amendment to the BSSN system as presented above has been the switch to *densitized lapse*. This quantity is defined as

$$Q \equiv e^{-6n\phi}\alpha. \tag{2.23}$$

If this replacement is made in the BSSN evolution equations, the stability is greatly improved. See, for instance, [123] for numerical examples where previously unstable evolutions achieved stability through this substitution. The precise mathematical reason for the advantages of densitizing the lapse is not known, but various authors have shown that equivalent conditions are necessary to achieve strong [87] or symmetric [72] hyperbolicity. Khokhlov & Novikov [86] are able to prove some desirable properties of general 3+1 formulations using gauges of this type. Nagy et al [97] have produced a simpler evolution system, based on the standard ADM equations, but with a densitized lapse; both this system and BSSN with the densitized lapse are hyperbolic. Another way to look at densitizing the lapse is that it is a cheap way of obtaining dynamic slicing in the system; the lapse function gains a time dependence through ϕ.

A remaining area of instability in fully 3D evolutions is encountered when we relax the symmetry assumptions that allow us to restrict the evolution domain to octant symmetry. Even runs of a single IEF black hole in quadrant symmetry were prone to an instability that would crash the simulation after a few hundred M. This seems to have been overcome, however, by use of a trick suggested by [5]: we add to the evolution equation for $\tilde{\Gamma}^i$ a term containing the algebraic constraint that defines $\tilde{\Gamma}^i$ (2.17):

$$\begin{aligned}\partial_t \tilde{\Gamma}^i &= \tilde{\gamma}^{jk}\partial_j\partial_k\beta^i + \frac{1}{3}\tilde{\gamma}^{ij}\partial_j\partial_k\beta^k + \beta^j\partial_j\tilde{\Gamma}^i - \tilde{\Gamma}^j\partial_j\beta^i \\ &\quad + \frac{2}{3}\tilde{\Gamma}^i\partial_j\beta^j - \left(\chi + \frac{2}{3}\right)\left(\tilde{\Gamma}^i - \tilde{\gamma}^{jk}\tilde{\Gamma}^i_{jk}\right)\partial_j\beta^j \\ &\quad - 2\tilde{A}^{ij}\partial_j\alpha + 2\alpha(\tilde{\Gamma}^i{}_{jk}\tilde{A}^{jk} + 6\tilde{A}^{ij}\partial_j\phi - \frac{2}{3}\tilde{\gamma}^{ij}\partial_j K). \end{aligned} \tag{2.24}$$

As this vanishes analytically, the physical evolution system is unchanged – the effect of this additional term is to alter the characteristic structure of the basic equations and the constraint subsystem. Yo et al [137] find that this trick aids in moving from octant symmetry to full-domain evolutions.

There are many other possible adjustments that can be made to the BSSN system that affect its stability properties. Yoneda & Shinkai [139] investigate a very large parameter space of such adjustments. Apart from the correction above, we have not applied any of these, although further adjustment may be required for highly spinning black holes.

2.3 Slicing Conditions

An issue that has been referred to only peripherally up to now is the question of *slicing conditions*[1] – the determination of the lapse and shift throughout the evolution.

The standard ADM system (Section 2.1) and the BSSN conformal traceless system (Section 2.2) are alternative mathematical expressions of the evolution of the physical fields γ_{ij} and K_{ij}. Neither system says anything about the lapse and shift as they are gauge quantities – manifestations of GR's insensitivity to coordinates. Nevertheless, we must make some statement about them at every stage of an evolution, and it seems that a poor choice of slicing can have a definite impact on the quality and stability of an evolution.

We can categorize slicing conditions in three or four ways. Conceptually the simplest thing to do is *analytical slicing* – supplying an explicit formula for α and β^i in terms of the space-time coordinates x^μ (and perhaps additional parameters, such as black hole masses). The problem with this approach is that we don't know what a good analytical choice might be. The simplest possible choice in general would be *geodesic slicing*: $\alpha = 1$, $\beta^i = 0$ everywhere. However, this is an extremely poor choice in almost all situations, as every point on the coordinate grid will fall into a black hole singularity in a finite time.

Another option for specifying α and β^i is *algebraic slicing* – writing the lapse and shift as algebraic combinations of the live evolution fields. This approach has the virtue of tying the behavior of the slicing variables to that of the physical fields, but will not generally have good stability properties.

The third option is *differential slicing* – demanding that α and β^i satisfy some differential equations. This is an attractive option for certain types of initial data. For instance, setting $\partial_t K = 0$ yields (from equation (2.21)) the *K-freezing* condition for α,

$$\triangle\alpha = \beta^i \partial_i K + \alpha K_{ij} K^{ij}.$$

A similar condition for β^i, the *$\tilde{\Gamma}$-freezing* condition, can be obtained by setting $\partial_t \tilde{\Gamma} = 0$. These conditions can be very useful for the evolution of conformal data (where K is often set to zero, yielding so-called *maximal slicing*). However, these conditions have two general disadvantages for the types of evolution I am considering here. The first disadvantage is inefficiency. The K-freezing and $\tilde{\Gamma}$-freezing conditions are elliptic equations that must be solved on a time slice, much as with the constraints; but these have to be solved on *every* time slice. This is usually a huge undertaking, and will dominate the evolution time.

With a few notable exceptions (see, for instance, [111]), then, direct elliptic gauges are not used. Instead, many elliptic gauges can be transformed into – or at least approximated by – parabolic or hyperbolic differential equations. For instance, the *K-freezing* condition

[1] These are sometimes called "gauge conditions". Also, the shift conditions are sometimes called "spatial gauge conditions".

above can be approximated by a *K-driver* condition:

$$\partial_t \alpha = -\alpha^2 f(\alpha)(K - K_0).$$

The idea here is that the system will quickly relax to a state of unchanging K and hence α will also freeze (assuming a reasonable initial α is taken, of course)[2]. In fact, this condition is used in many modern codes, including MAYA with the particular function choice $f(\alpha) = 2/\alpha$:

$$\partial_t \alpha = -2\alpha(K - K_0). \tag{2.25}$$

This slicing is known as "1 + log", as for $K_0 = 0$ and $\beta^i = 0$, it has an analytic solution $\alpha = 1 + \log|\gamma|$. In MAYA it is implemented using the evolution equation for ϕ (2.19):

$$\partial_t \alpha = D_m \beta^m - \alpha K \equiv 6\partial_t \phi. \tag{2.26}$$

Additionally, since we often use the densitized lapse Q instead of α as a variable, we may derive from (2.26) an equivalent expression for Q's evolution:

$$\begin{aligned} \partial_t Q &\equiv \partial_t (e^{-6n\phi}\alpha) \\ &= e^{-6n\phi}\partial_t \alpha - 6n\alpha e^{-6n\phi}\partial_t \phi \\ &= 6e^{-6n\phi}(1 - n\alpha)\partial_t \phi \\ &= 6Q\left(\frac{1}{\alpha} - n\right)\partial_t \phi. \end{aligned} \tag{2.27}$$

A similar driver condition exists for the shift vector β^i. Based on the $\tilde{\Gamma}$-*freezing* elliptic equation, the *Gamma driver* condition is a parabolic equation:

$$\partial_t \beta^i = \lambda \partial_t \tilde{\Gamma}^i, \tag{2.28}$$

where λ is a positive real free parameter, which we have taken as $\lambda = 0.05$.

These driver evolution equations are much more efficient than the elliptic equations on which they're based, and generally enhance the stability of the evolutions in which they are used. However, they are still aimed at stationary, or near-stationary space-times. They have proved well suited to the stable evolutions of single non-moving black holes. Groups working on binaries on quasi-stationary circular orbits have been experimenting with slicings that actually freeze the motion of the binary in a co-rotating frame [36]. In this context, driver conditions may work, and even help enhance the stability of the orbits.

In more general physical situations, such as black-hole mergers or more general orbits, we may be forced to consider the actual movement of black holes across the numerical grid.

[2] This is a slight generalization of the Bona-Massó slicing conditions [26].

In these cases, driver conditions are not appropriate, as they try to achieve or maintain stationarity in the coordinates chosen, while the space-time is manifestly *not* stationary. I am unaware of any proposal for a dynamic evolution of α and β^i for situations where black holes are moving through the numerical grid; for the time being, we are forced to use analytic or algebraic gauge conditions for the parts of the evolution featuring moving singularities.

2.4 The PSU Maya Numerical Code

3D numerical work done for this thesis was performed using the Penn State Maya code (for an overview of the Maya code, see [118]). Maya was designed and built as a successor to Penn State's Agave code (itself a descendant of the Binary Black Hole Grand Challenge Alliance code). The chief architect of Maya was Erik Schnetter; since then it has seen significant changes at the hands of Fiske, Laguna, Shoemaker, Smith, Sperhake, and myself.

Maya is a unigrid Cartesian code that implements the BSSN system of equations in a method-of-lines evolution. The time integration algorithm used is iterated Crank-Nicholson (ICN). Spatial derivatives are calculated using centered-difference formulas[3] such as

$$\begin{aligned}
\frac{\partial u}{\partial x} &= \frac{u(x+\Delta,y,z)-u(x-\Delta,y,z)}{2\Delta}+O(\Delta^2) \\
&= \frac{u_{i+1,j,k}-u_{i-1,j,k}}{2\Delta}+O(\Delta^2), \\
\frac{\partial^2 u}{\partial x \partial y} &= \frac{1}{2\Delta^2}\left[u(x+\Delta,y+\Delta,z)+u(x-\Delta,y-\Delta,z)+2u(x,y,z)\right. \\
&\quad \left.-u(x+\Delta,y,z)-u(x-\Delta,y,z)-u(x,y+\Delta,z)-u(x,y-\Delta,z)\right] \\
&\quad +O(\Delta^2) \\
&= \frac{1}{2\Delta^2}\left[u_{i+1,j+1,k}+u_{i-1,j-1,k}x+2u_{i,j,k}\right. \\
&\quad \left.-u_{i+1,j,k}-u_{i-1,j,k}-u_{i,j+1,k}-u_{i,j-1,k}\right]+O(\Delta^2).
\end{aligned}$$

The meaning of the $O(\Delta^2)$ terms above is that the error in all physical quantities should scale with the square of the spatial grid spacing: Maya is inherently a *second-order-accurate* code. As an example, Fig 2.3 shows a point-wise comparison of the error in the Hamiltonian constraint C_H for an IEF single-BH space-time. Analytically, C_H should be zero. The actual error seen is second-order, as the spatial derivatives used to calculate the constraint have $O(\Delta^2)$ errors. It is important to bear in mind that second-order convergence does not mean

[3] The exceptions to this are *advection terms*, of the form $\beta^k \partial_k u$; here one-sided derivatives are used, though the derivative accuracy is maintained – see, for instance [83], eqn (3.18).

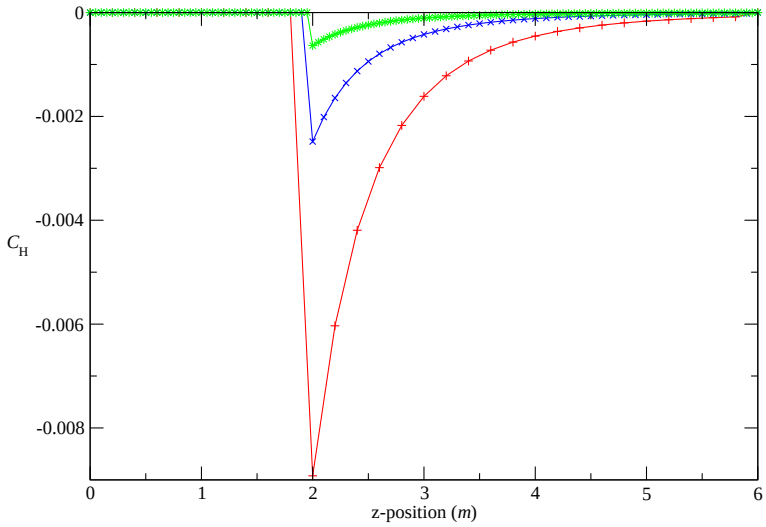

Figure 2.3: Constraint violation of IEF data for three different resolutions: $\Delta = 0.2m$ (+), $0.1m$ ($\times$), and $0.05m$ ($*$).

that the errors encountered here will actually be smaller than those of a first-order-accurate code, nor that they will be larger than those of a fourth-order-accurate one. The coefficient in front of Δ^2 in error terms marked $O(\Delta^2)$ could be very large, or very small. When dealing with black holes, most non-constant fields vary strongly with inverse powers of r near the singularity (see, for instance, the standard ADM (A.5) or BSSN (A.6) fields for the IEF solution). Such a falloff results in steep gradients in the fields and spatial derivatives at all orders; the errors our code gives rise to will also exhibit these strong gradients.

To investigate the order of convergence, we often calculate the *convergence ratio*, C. This comes in two flavors, two-level and three-level. The two-level ratio assumes we know the analytic (exact) value of a field, and compares the errors at two resolutions:

$$C_2(u) \equiv \frac{u_{2\Delta} - u_0}{u_\Delta - u_0}. \tag{2.29}$$

The three-level ratio assumes we know nothing of the analytic value of the field; it compares function values at three different resolutions:

$$C_3(u) \equiv \frac{u_{4\Delta} - u_{2\Delta}}{u_{2\Delta} - u_\Delta}. \tag{2.30}$$

Then nth-order convergence is indicated by $C_2 = 2^n$ (or $C_3 = 2^n$).

It should be stressed that *convergence does not mean correctness*. For instance, initial data that analytically violates the constraints will still be convergent, just not to zero.

The MAYA code is itself built upon the CACTUS [1] infrastructure. CACTUS is an open code designed with numerical physics evolution in mind. It automatically handles the parallelization of execution over multiple processors. The CACTUS code base also contains extensive GR-specific code for NR; however, for the most part this is ignored in favor of locally written MAYA code; we have maintained independence from the CactusEinstein code for several reasons, not the least of which is our belief that we can supply a simpler and cleaner Einsteinian evolution system than is supplied in CACTUS. Nevertheless, we have made use of certain CACTUS GR utilities, including the CACTUS AHFinder thorn for locating apparent horizons (see Section 3.5).

MAYA's chief advantage over previous codes in this line is its simplicity. It is primarily written in Fortran 90. The only formulations coded in are the BSSN system, with slight variations. The evolution scheme uses the method of lines (MOL) with an iterated Crank-Nicholson (ICN) time integrator[4]. Simple choices like these have led to a very readable, and easy-to-maintain evolution code. While some of the limitations mentioned make it unlikely that the code as it is will be able to generate production-level wave forms for gravitational-wave physicists, it is an ideal testing ground for several new ideas in numerical simulations – for example, singularity excision (see Section 2.7 below).

MAYA has been a participating code in the "Apples with Apples" NR collaborative effort [2].

2.5 Kerr-Schild Coordinates

The Kerr solution is the unique solution to Einstein's equations for a vacuum space-time exhibiting stationarity and asymptotic flatness. That is to say, any isolated, spinning black hole must be a Kerr black hole. Such a space-time is parametrized by exactly two quantities: the black hole mass M and the spin angular momentum, $\vec{J}$. We usually deal instead with the angular momentum per unit mass, $\vec{a}$. A feature of this solution is that the spin magnitude $a \leq M$. A solution with $a = M$ is called "extremal". It is believed that any star undergoing collapse to a spinning black hole will be forced to boil off excess angular momentum until it can satisfy this condition.

Originally, the Kerr solution was written in Boyer-Lindquist coordinates (see Appendix Section A.2.1). These reduce in the zero-spin case to the Schwarzschild solution in Schwarzschild coordinates. In particular, the event horizon coincides with a coordinate singularity.

The MAYA code is unusual amongst modern evolution codes in that it works almost

[4]Other time-integration algorithms, such as Runge-Kutta, are possible, but have not been implemented to date.

exclusively with black holes expressed in Kerr-Schild coordinates. These coordinates are *horizon-penetrating* – there is no coordinate singularity at the location of the horizon, as there are in Schwarzschild or Isotropic Schwarzschild coordinates. This horizon-penetrating property means that we are free to include points within the horizon in our evolution (we still need to be careful with the physical singularity, of course). For this reason, Kerr-Schild data is ideal for singularity excision, a central feature of the Maya code.

Another attractive feature of Kerr-Schild coordinates is that it allows a closed-form solution for single black holes with spin angular momentum (that is, the Kerr solution), and also allows us to apply boosts transparently (see Appendix Section A.3). Thus the generic final state of an arbitrary black hole interaction – a massive, spinning black hole with net linear momentum – is realizable in Kerr-Schild coordinates. This is in stark contrast to well-known conformal data schemes such as Bowen-York (see Appendix Section A.4.2), where spin and linear momentum come at the cost of violating the Hamiltonian constraint.

2.6 Outer Boundaries

Numerical Relativity, as approached by Maya and most other codes, is an initial-boundary-value problem. That is, having determined the appropriate initial data to begin with, we must evolve the spatial metric and extrinsic curvature "in the bulk" and also worry about how the fields behave on the inner and outer boundaries. The fields cannot be evolved with the standard ADM or BSSN equations, as the time derivatives given by (2.18) – (2.22) require first- and second-order spatial derivatives, which are calculated using centered difference formulas. Generically, the boundaries we must deal with can be split into two types, which are dealt with in very different ways.

The simplest possible boundary condition is to do nothing at all, that is, to "freeze" the fields on the boundary so that their value remains what they started with on the initial slice. This we refer to as the *frozen* boundary condition. Freezing fields like this is conceptually acceptable if there is no dynamics in play: if we are dealing with a stationary situation (for example, flat space in Minkowski coordinates, or a single Schwarzschild black hole in IEF coordinates), then there is no physical need for these fields to change at all. For a more complicated problem with a time-dependent solution, this cannot be correct.

The next simplest condition is analytically enforcing the form of the fields on the boundary at each step – *Dirichlet* conditions. In contrast to the frozen conditions, this at least allows us a time-dependent solution. However, in practice we are still very unlikely to know the form of the fields on the boundaries if we don't know them in the bulk, so Dirichlet conditions should not be expected to fare much better than frozen.

Similar to Dirichlet in this scenario are *Neumann* conditions, where instead of enforcing the form of the fields themselves, we enforce some of the spatial derivatives. Though this

distinction is important in initial-value problems, it is not in evolution problems. It still requires more-or-less perfect knowledge of the dynamical fields on the boundary.

In physically interesting evolutions, therefore, we need approaches to boundaries that account for our imperfect prior knowledge of the space-time. The "correct" thing to do from a mathematical point of view is only now beginning to be understood; much work has been done on mathematically self-consistent and well-posed boundary conditions for NR problems ([37], [38], [70]). Most of the proposed algorithms, however, are tied to a particular evolution system with special properties, properties that might not be possessed by the BSSN system. Moreover, most proposals are tested only for linearized systems.

In lieu of useful guidance from mathematicians, we use more pragmatic approaches for both inner and outer boundaries in the Maya code. Schemes that allow for more ignorance on the part of the numericist include *extrapolation* schemes and *radiative* schemes.

2.6.1 Radiative Outer Boundary Conditions

It is an unfortunate fact that a unigrid code, even with computational resources vastly superior to those available five years ago, simply cannot handle outer boundaries in an appropriate manner, without resorting to some kind of warped coordinate system (for example, so-called "fish-eye" coordinates [17]). In terms of the the black-hole problems we are interested in, current resolutions and computer power still leave us, at best, in what Thorne calls the "weak-field near zone" [132]. In this region, radiation does not readily resemble a linear perturbation of flat space. The problems we are trying to solve have natural boundary conditions at spatial (or null) infinity; nothing obvious applies at the physical distances we have to deal with. It is natural to expect therefore, that whatever wave forms we may extract will not be serious candidates for detector observations.

Radiative boundary conditions assume that we know something about the solution near the boundary – perhaps the monopole or dominant fall-off behavior. Then we can treat the deviation of the numerically evolved data near the boundary from this background behavior as an error that can be radiated away. This then gives rise to an evolution equation for the field itself:

$$\begin{aligned} u &= u_0 + \frac{w(r-vt)}{r^n} \\ \Rightarrow \partial_t\, u &= -v\frac{x^i}{r}\partial_i(u-u_0) - v\frac{n(u-u_0)}{r}, \end{aligned} \tag{2.31}$$

assuming the background u_0 has no time dependence. This scheme should work well under a few assumptions:

(i) We know the background field parameters.

(ii) The physical field really is close to the background field, so that the deviations are in the perturbative regime.

(iii) The deviations for each field propagate with a single well-defined velocity v.

These conditions can easily be satisfied for certain perturbative situations, for instance weak Teukolsky or Brill waves on a static Schwarzschild background. The weakness guarantees that (i – ii) hold, and the fact that the dominant errors will be gravitational waves, which have a well-defined coordinate speed on a weak background, guarantees that (iii) is satisfied.

This should also work for single-BH evolutions of the sort that have been done stably in 3D only since the mid-1990s. In these cases, the analytic solution was well known, and could be used as the background for these outgoing boundary conditions. Thus the "radiation" that is being transported out of the grid is simply the numerical differencing error inside the bulk.

In fact, all simulations of interest to us in vacuum general relativity should satisfy (iii) for the physical fields. However, not all modes propagating in the domain are physical (gravitational wave) modes: there are also gauge modes and constraint-violating modes, into which we have much less physical insight. Moreover, as the BSSN evolution variables are not characteristic variables, each BSSN field will contain a mixture of characteristic modes, with possibly distinct characteristic velocities. Obviously multiple velocities will violate (iii).

If, however, we do not know know what the parameters are – either because of the manner of construction of the initial data, or because of the strength of the perturbation – then we are at something of a loss. For the case where we expect a single BH to emerge as the end result of the evolution, we may use global quantities such as the ADM energy as the final mass of the hole. This approach has two problems: first, to calculate the ADM energy in a space-time, we need to be able to integrate over a sphere at infinity – a finite-radius integration will be a poor estimator; second, assuming we know the ADM energy perfectly, this will still definitely overshoot the actual BH mass, as we know that radiation will be emitted to infinity, and will take energy with it. To enter the appropriately reduced mass, we would need to know how much radiation would be emitted, and this is a quantity we can only determine from the evolution to begin with.

2.7 Inner Boundaries and Singularity Excision

One of the characteristics of general relativistic problems that does not seem to appear in other classical field theories (like say, electrodynamics), is the presence of real physical singularities that represent regions of infinite curvature. Interesting as these singularities are,

their presence causes immediate and drastic problems for any computer-based investigation: a computer simply cannot represent infinity, and the attempt will inevitably crash or at the least pollute the simulation. Therefore, we need an approach to minimizing the effects of singularities. For a comprehensive review of this topic, see Kenneth Smith's thesis [121].

2.7.1 Singularity-Avoiding Slicing

Two general approaches are in widespread use today. The first is *singularity-avoiding slicing*. This approach assumes that the initial data slice does not contain a singularity, and slows down the march of coordinate time in regions of increasingly steep curvature, to the point that evolution is essentially "frozen" near a black hole. It achieves this by forcing the collapse of the lapse function. Though an effective method for short periods of time, the increasingly steep gradients encountered in the lapse as it collapses inevitably lead to the failure of the evolution.

The singularity-avoidance technique is well adapted to certain types of data, such as the Misner, Brill-Lindquist, and Bowen-York type. In these cases, the use of throats and many-sheeted topologies means that singularities are not present on the initial slice. In general, the "puncture" evolutions of Brügmann et al [30, 34, 35] use this singularity avoidance procedure.

2.7.2 Singularity Excision

The other prominent approach in NR is a technique known as *singularity excision*. The idea, usually attributed to Unruh[5] is a conceptually simple one: we know that the event horizon of a black hole is causally disconnected from the outside universe. Therefore we should be able simply to remove – excise – the singularity from the computational domain without affecting the physics outside the horizon.

This technique requires, however, that we have a coordinate scheme that is well behaved everywhere up to the excision boundary, which must be inside the horizon. Schwarzschild coordinates (Section A.1.1) fail in this regard, as do Isotropic Schwarzschild (Section A.1.2), but Kerr-Schild-type coordinates are horizon-penetrating – they are well behaved all the way to the physical singularity. The black hole data used in Maya tends to be of the Kerr-Schild type, especially the zero-spin Ingoing Eddington-Finkelstein coordinates (Section A.1.3).

The Maya code can excise regions of space around a singularity using one of three shapes: cubical, spherical, and cuboctahedral. In each case, the parameters that determine the final shape and size of the excision region (or "mask") are the grid spacing and the *excision radius*. The latter determines an ideal spherical surface into which the mask must be inscribed. Fig 2.4 shows part of cubical and spherical masks for the same excision radius

[5] This is quoted as a private communication with Jonathan Thornburg [130].

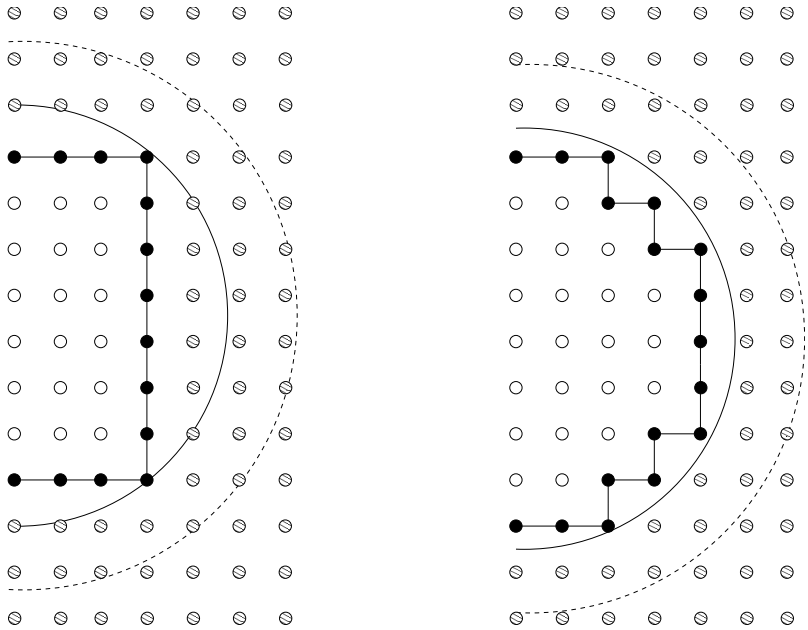

Figure 2.4: Excision mask for cubical and spherical excision at the same resolution. The dashed curve is the location of the apparent horizon, while the solid curve is the excision radius. Points are colored according to their status: excised (white), excision boundary (black), and evolution (shaded).

and the same grid resolution. Obviously, the quality of the approximation to a Platonic ideal sphere increases with resolution.

The original approach to excision was to ignore totally all excised points (including those on the boundary of the region): when a spatial derivative term had to be calculated at a nearby point, a specialized stencil was formed – a set of neighboring live points with weights that would give the desired derivative to the appropriate order, without having to use excised points. For a complex excision mask shape, many different templates would be needed to cover all possibilities. This method was used in the earliest 3D implementations of black holes, especially those carried out by the Binary Black Hole Grand Challenge Alliance [48].

A major breakthrough came in [5], with what is called "simple excision". The idea was to first extrapolate data onto the excision boundary, and then use normal centered differencing at all live grid points. The centered differencing is safe everywhere, assuming the extrapolation onto the boundary is of high enough order. Initially this was done only on a cubical boundary (the presence of edges and corners meant that this was still not trivial), but the Maya group have successfully applied it to the other shapes mentioned above. Moreover, the "simplest" part of simple excision – the copying over of time derivatives (zeroth-order extrapolation) from the nearest interior point to the boundary point – is

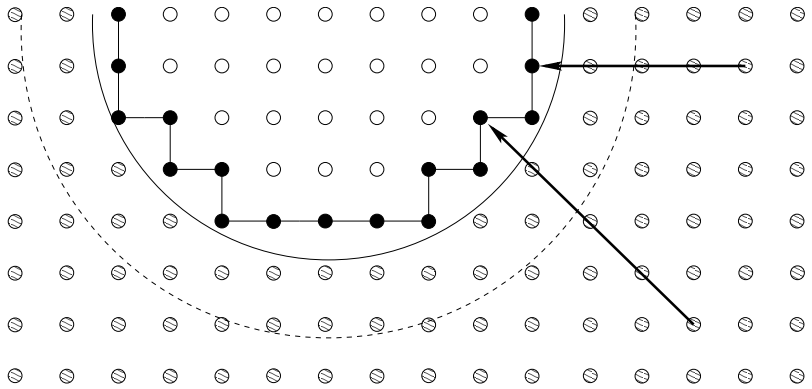

Figure 2.5: Extrapolation onto excision boundary points uses the closest on-grid normal direction, and requires data from up to four live points along that normal.

sensible only for actual stationary evolutions. The MAYA excision routines use up to four interior points along an extrapolation direction to fill in data on the boundary. The arrows in Fig 2.5 indicate extrapolation directions for certain points on the cubical excision boundary.

There are two ways to use extrapolation to update the data on an excision boundary point: evolve data on all live points and extrapolate the result onto the boundary (*solution extrapolation*), or extrapolate the time derivatives onto the boundary, and then use this to evolve to the next time level (*right-hand-side extrapolation*). In general, both schemes work, though right-hand-side extrapolation does seem to enhance stability in stationary situations.

2.7.3 Moving Holes and Dynamic Excision

The foregoing is sufficient to evolve stationary space-times such as a single Kerr-Schild black hole. A more rigorous test is the ability to move an excision mask across the numerical grid. Such *dynamic excision* is the hallmark of the MAYA code; see [118, 121, 123] for recent developments. As the hole moves across the grid, grid points in its path become part of its excision mask, while other already excised points find themselves "live" again. It is easy to excise points from the domain – we simply zero them out. However, the newly revealed points have to have meaningful data assigned to them. This procedure is called *repopulation*, and is achieved roughly as follows: each of these points finds itself now on the excision boundary. Every point on the boundary has a well-defined normal direction, as used for the normal updating scheme. Now these directions are used to extrapolate data straight onto the boundary point in the same way that the solution or time derivative would be extrapolated to during evolution. As with the solution and right-hand-side extrapolation, the population procedure has an order; by default we use a second-order extrapolation scheme.

The first tests MAYA was put through with dynamic excision involved a black hole moving in a plane of the domain, with motion entirely due to a coordinate transformation. These were not tests of any realistic physical scenario, but of the ability of the code to allow the singularity to move about the grid without poking out of the excision region, or triggering or exacerbating unstable modes.

The "circling" black hole runs [118] used a prescribed coordinate motion, which yielded the position of the center of the hole for the purposes of calculating the lapse and shift, and also for centering the excision mask.

2.7.4 Singularity Tracking and the GaussTracker Module

Until recently, the same data (an analytic trajectory) was used to update both the singularity position and the excision mask center in dynamic evolutions. This is troubling in principle: the idea of a Cauchy evolution is that the initial data (plus some sensible boundary scheme) should be sufficient to determine the future history of the initial slice.

A recent improvement of this situation in the MAYA code has been the introduction of GaussTracker, a mask update algorithm that allows us to *not* feed the hole position back into the mask. This algorithm looks at the shape of one of the evolution variables – typically K – on the current time step, along axes passing through the center of the current excision mask. For each axis direction, a gaussian fit is used to assess the center of curvature of the variable. When this has been done for all three directions, the mask is re-centered to lie as close as possible to the new "ideal" position.

This algorithm has been rigorously tested in Sperhake et al [123], with a single Kerr-Schild black hole performing a coordinate in-spiral and out-spiral in the numerical domain. The runs are very stable, going for thousands of M with performance rivaling that of the trajectory-determined excision centering.

The availability of such a technique may prove essential in dealing with situations where the motion is not purely coordinate-based. I will address this further in Chapter 5.

Chapter 3

Initial Binary Data and Diagnostics

To be useful and relevant to astrophysics, General Relativity must be able to describe interactions that we might actually observe. If we restrict ourselves to non-matter situations, the most interesting scenario is the general interaction of two black holes that fall into mutual orbit and eventually merge. Such an interaction should produce a continuous and highly characteristic gravitational wave-train, observable (if sufficiently powerful) by our ground-based and future space-based GW detectors (e.g., LIGO, GEO, VIRGO, LISA).

Unfortunately, no analytic solutions exist for the entire history of a binary black hole system. In fact, even finding initial data that satisfies the constraints is not an easy task. Misner [93, 94] and Brill & Lindquist [33] developed slightly different schemes that at least satisfied the constraints initially.

The desirability of such information was appreciated long before we could actually envisage a way to produce it. The problem as a whole has generally been split up into three distinct realms: the far limit (FL), where the holes are sufficiently far apart that approximation methods can be used to extract physical wave forms; the close limit (CL), where a different approximation scheme can supply full wave forms for the settling down of the final black hole; the full numerics (FN) region, where no obvious approximation scheme will be sufficient.

The FL region has been treated very successfully using varying levels of post-Newtonian (PN) expansions (for a review see, e.g., [25]). Perturbative studies of the CL region [3, 106] have supplied late-merger and ring-down wave forms also. What remains, when both of these schemes have been pushed past the limits of their validity, is the strong-interaction region, the domain of full NR.

While 3D codes were being developed, more precise knowledge was developed about the close-limit behavior of both Misner and Brill-Lindquist data [3, 106]. Such calculations gave very precise tests of the fully numerical 2D evolutions of these data sets, as well as supplying some insight into the radiation content of the different sets [3].

The binary-black hole head-on collision was one of the first multiple-hole problems to be addressed. Its axisymmetry made full numerical evolution at least plausible in the late 1960s and early 1970s. In the next three chapters, I will present recent work on the head-on collision problem that will indicate how much progress has been made in this area. This is

original work, carried out in collaboration with Pablo Laguna, Kenneth Smith, and Ulrich Sperhake, and has not been published at time of writing.

This chapter is laid out as follows: in sections 3.1, 3.2 and 3.3, I give an introduction to the initial data used in our simulations, and the numerical method used to solve the elliptic constraint equations for this data. In Section 3.4, I discuss the related issue of an appropriate analytic slicing condition for this data. In sections 3.5, 3.6, 3.7 and 3.8, I introduce the diagnostic measures used in analysis of the evolving space-time.

3.1 Initial Binary BH Data

The data we use in MAYA for binary evolutions is the so-called "HuMaSh" data, first defined in [92], and used in binary grazing-collision evolutions in [31]. This data uses a very simple scheme, based on simply adding Kerr-Schild solutions in an obvious way:

$$\begin{aligned} \gamma_{ij} &= {}^{A}\gamma_{ij} + {}^{B}\gamma_{ij} - \eta_{ij}, \\ K_{ij} &= \gamma_{(i|k} \left[{}^{A}K^{k}{}_{|j)} + {}^{B}K^{k}{}_{|j)} \right]. \end{aligned} \tag{3.1}$$

For example, in the case of zero boost and zero spin for either hole, the three-metric would be (in Cartesian coordinates):

$$\gamma_{ij} = \delta_{ij} + \frac{2\, m_A\, x_{Ai}\, x_{Aj}}{r_A^3} + \frac{2\, m_B\, x_{Bi}\, x_{Bj}}{r_B^3},$$

where m_A and m_B are the bare masses of the pre-merger holes, and $r_A \equiv \sqrt{\delta^{ij}\, x_{Ai}\, x_{Aj}}$ is the coordinate distance from singularity A. This initial data was suggested largely because of the experience already gained in singularity excision with Agave, the precursor to the MAYA code. The chief attraction of a Kerr-Schild superposition is the singularity-penetrating nature of the coordinates. As noted above, individual Kerr-Schild black holes possess spin, and can be boosted analytically. For simulations involving initially well-separated black holes, these properties arguably make the data a better approximation to real astrophysical objects than Bowen-York multiple-hole data.

The HuMaSh data automatically obeys the obvious limits one would expect: as the mass of one hole shrinks to zero, it tends to the exact Kerr-Schild solution for a single hole; in the far-field limit, it tends to an exact Kerr-Schild single-hole solution for a hole of mass $M = m_A + m_B$. Nevertheless, it is manifestly not going to satisfy the constraints *ab initio*; Fig 3.1 shows the Hamiltonian constraint along the z-axis for three different spatial resolutions. The central differencing used in spatial derivatives breaks down badly on the excision boundary (left frame), but even safely removed from the boundary, the constraint

converges to a non-zero form (right frame).

Marronetti et al have suggested a way of attenuating this data with an exponential damping factor, to lessen the violations near the singularities [91]. While considering this attenuated data, Bonning et al [27] have provided a fairly detailed investigation of the general properties of the underlying HuMaSh data. In particular, they show that it has the correct gravitational binding energy in the Newtonian far-separation limit.

I should note here that the HuMaSh prescription, and variants such as [91], are not the only conceivable way of going about using the Kerr-Schild solution to generate binary data. An alternative method involving an initial superposition of the Kerr-Schild four-metric was suggested by Bishop et al [23, 24]. Though this data captures the spirit of Kerr-Schild data, and has some nice perturbative properties, the constraints are generally quite complex, and not always elliptic.

Other Kerr-Schild-based schemes have been developed by Moreno et al [96] and Dain [53, 54]. Apart from a perturbative evolution of the latter data [109], I am unaware of simulations using these data sets.

3.2 Constraint Decompositions

As it does not satisfy the constraints, we must condition the HuMaSh data using an elliptic solver. The constraint equations (2.9) and (2.10) are a set of four coupled PDEs involving the three-metric and the extrinsic curvature. To solve them, we must first write them in terms of a set of four functions, which can be solved for. There is, obviously, an infinite number of ways in which to do this. The most direct – solving for four independent components of K_{ij}, for instance – have the benefit of simplicity, but will not respect the symmetries of the physical problem, and will quite probably lead to an ill-posed system.

3.2.1 The Conformal Transverse-Traceless Decomposition

York [141] outlined a method of decomposing the ADM system using a conformal factor ψ, which I review briefly here. We begin by isolating the trace-free part of K_{ij}:

$$A^{ij} = K^{ij} - \frac{1}{3}\gamma^{ij}K.$$

With this split, the constraints become:

$$C_H \equiv R + \frac{2}{3}K^2 - A_{ij}A^{ij}, \tag{3.2}$$

$$C^i_M \equiv \nabla_j A^{ij} - \frac{2}{3}\nabla^i K. \tag{3.3}$$

Next, we choose a conformal rescaling of the metric and the trace-free part of the

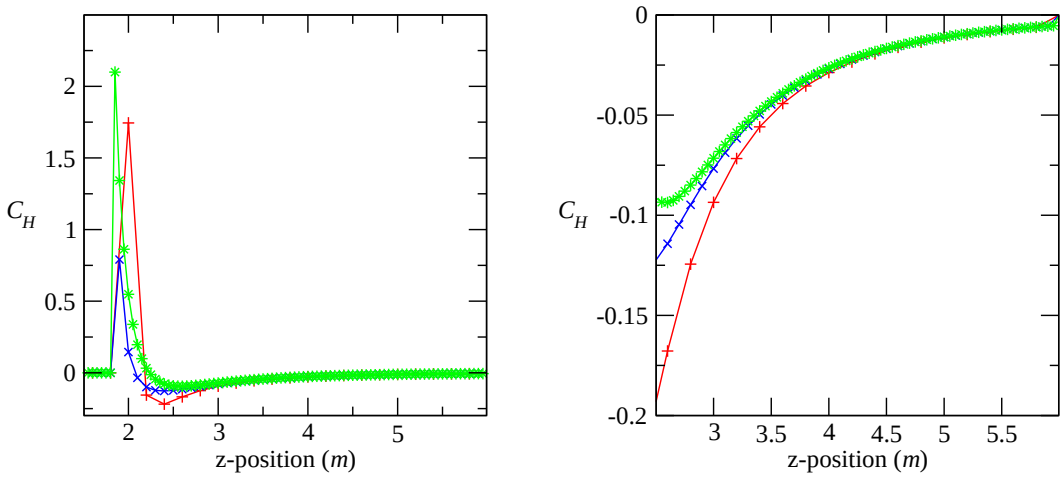

Figure 3.1: Hamiltonian constraint violation C_H outside unsolved small-separation data for different resolutions, coarse ($+$), medium ($\times$) and fine ($*$). (i) The violation along z-axis, including excision boundary. (ii) Zoom-in on section of z-axis away from excision region.

extrinsic curvature:

$$\begin{aligned} \gamma_{ij} &= \psi^4 \hat{\gamma}_{ij} \Leftrightarrow \gamma^{ij} = \psi^{-4} \hat{\gamma}^{ij}, \\ A_{ij} &= \psi^{-2} \hat{A}_{ij} \Leftrightarrow A^{ij} = \psi^{-10} \hat{A}^{ij}. \end{aligned}$$

[1] One can choose K to have no conformal scaling, so that $K = \hat{K}$. With these transformations, the constraint equations become

$$\begin{aligned} C_H &\equiv -8\psi^{-5} \hat{\triangle}\psi + \psi^{-4}\hat{R} + \frac{2}{3}K^2 - \psi^{-12}\hat{A}_{ij}\hat{A}^{ij}, \\ C^i_M &\equiv \psi^{-10}\hat{\nabla}_j \hat{A}^{ij} - \frac{2}{3}\psi^{-4}\hat{\nabla}^i K. \end{aligned}$$

The final step in this decomposition is to note that in the momentum constraint, $\hat{A}^{ij}$ appears only in its divergence, so that its transverse part $\hat{A}^{ij}_T$ is not restricted by the equations and can thus be chosen freely. We can express the longitudinal part, $\hat{A}^{ij}_L$, as the symmetric, traceless gradient of some vector field X^i:

$$\begin{aligned} \hat{A}^{ij}_L &= \hat{\nabla}^i X^j + \hat{\nabla}^j X^i - \frac{2}{3}\hat{\gamma}^{ij}\hat{\nabla}_k X^k \equiv (\hat{L}X)^{ij}, \\ \Rightarrow \hat{\nabla}_j \hat{A}^{ji} &= \hat{\nabla}_j \hat{\nabla}^j X^i + \frac{1}{3}\hat{\nabla}^i \hat{\nabla}_j X^j + \hat{R}_j{}^i X^j \equiv (\hat{\triangle}_L X)^i. \end{aligned}$$

Now the momentum constraint can be written

$$C^i_M = \psi^{-10}(\hat{\triangle}_L X)^i - \frac{2}{3}\psi^{-4}\hat{\nabla}^i K. \tag{3.4}$$

The advantage of this decomposition is that it isolates the degrees of freedom in the problem. Initially there were twelve quantities, the components of the (symmetric) three-metric and extrinsic curvature tensors. There are now four quantities that need to be computed when solving the four constraint equations: the conformal factor ψ and the components of the vector X^i. The conformal metric $\hat{\gamma}_{ij}$, the trace of the extrinsic curvature K, and the transverse conformal trace-free part of the extrinsic curvature $\hat{A}^{ij}_T$ can be specified freely.

However, this last is a non-trivial task. It would be far easier to specify instead some tensor that is merely traceless. This can be achieved as follows: take any symmetric traceless tensor $\hat{M}^{ij}$. This then can be decomposed into transverse and longitudinal parts:

$$\hat{M}^{ij} \equiv \hat{A}^{ij}_T + (\hat{L}Y)^{ij}.$$

[1] I use carets throughout this section to denote quantities in conformal space.

Then the decomposition of our $\hat{A}^{ij}$ can be written as

$$\hat{A}^{ij} = (\hat{L}V)^{ij} + \hat{M}^{ij},$$

where $V^i \equiv X^i - Y^i$, and $\hat{M}^{ij}$ can be specified as *traceless*, but not necessarily *transverse*. The price we pay for this freedom is a slightly more complicated momentum constraint, while altering the Hamiltonian constraint:

$$C_H = -8\psi^{-5}\hat{\triangle}\psi + \psi^{-4}\hat{R} + \frac{2}{3}K^2 - \psi^{-12}\hat{A}_{ij}\hat{A}^{ij}, \tag{3.5}$$

$$C^i_M = \psi^{-10}(\hat{\triangle}_L V)^i - \frac{2}{3}\psi^{-4}\hat{\nabla}^i K + \psi^{-10}\hat{\nabla}_j \hat{M}^{ji}. \tag{3.6}$$

This is the final *Conformal Transverse-Traceless* (CTT) system. It is, in fact, the system used to solve the initial data in all our 3D runs. Before proceeding, however, I will mention two alternative decompositions, also popular for initial-data calculations.

3.2.2 The Physical Transverse-Traceless Decomposition

An alternative to the above decomposition, with its own advantages, is the *Physical Transverse-Traceless* decomposition. It starts out in the same way, with the definition of the conformal factor and metric. However, now the traceless part of the extrinsic curvature is decomposed into transverse and longitudinal parts *before* conformal scaling is imposed:

$$A^{ij} \equiv (LW)^{ij} + A^{ij}_T,$$

and this A^{ij} is subject to the same conformal scaling as the three-metric:

$$\begin{aligned} \gamma_{ij} &= \psi^4\hat{\gamma}_{ij} \Leftrightarrow \gamma^{ij} = \psi^{-4}\hat{\gamma}^{ij}, \\ A_{ij} &= \psi^4\hat{A}_{ij} \Leftrightarrow A^{ij} = \psi^{-4}\hat{A}^{ij}. \end{aligned}$$

Then the Hamiltonian and momentum constraints become:

$$C_H = -8\psi^{-5}\hat{\triangle}\psi + \psi^{-4}\hat{R} + \frac{2}{3}K^2 - \hat{A}_{ij}\hat{A}^{ij}, \tag{3.7}$$

$$C^i_M = \psi^{-4}(\hat{\triangle}_L W)^i + 6\psi^{-5}(\hat{L}W)^{ij}\hat{\nabla}_j\psi - \frac{2}{3}\psi^{-4}\hat{\nabla}^i K. \tag{3.8}$$

As with the CTT decomposition, we avoid the awkwardness of finding a transverse A^{ij}_T by defining it as the transverse part of some generic symmetric traceless tensor $\hat{M}^{ij}$:

$$A^{ij}_T \equiv \psi^{-10}\hat{M}^{ij} - (LZ)^{ij}.$$

Then, defining $V^i \equiv W^i - Z^i$, the momentum constraint becomes

$$C^i_M = (\hat{\triangle}_L V)^i + 6\psi^{-1}(\hat{L}V)^{ij}\hat{\nabla}_j\psi - \frac{2}{3}\hat{\nabla}^i K + \psi^{-6}\hat{\nabla}_j\hat{M}^{ij} = 0. \tag{3.9}$$

This, together with the Hamiltonian constraint above, form the *Physical Transverse-Traceless* (PTT) system.

3.2.3 The Conformal Thin-Sandwich Decomposition

A drawback of the previous two schemes is their arbitrariness – there is no obvious connection to the dynamics. If, however, we use the definition of the extrinsic curvature in terms of the evolution of the three-metric, we can involve the lapse α and shift vector β^i in the decomposition:

$$\begin{aligned}
\alpha &= \psi^6\hat{\alpha}, \\
\gamma_{ij} &= \psi^4\hat{\gamma}_{ij}, \\
A^{ij} &= \psi^{-10}\hat{A}^{ij}, \\
\hat{A}^{ij} &= \frac{1}{2\hat{\alpha}}\left[\left(\hat{L}\beta\right)^{ij} - \hat{u}^{ij}\right],
\end{aligned}$$

leading to the following form for the constraints:

$$C_H = -8\psi^{-5}\hat{\triangle}\psi + \psi^{-4}\hat{R} + \frac{2}{3}K^2 - \psi^{-12}\hat{A}_{ij}\hat{A}^{ij}, \tag{3.10}$$

$$\begin{aligned}
C^i_M = {} & \frac{1}{2\hat{\alpha}}\psi^{-10}\left(\hat{\triangle}_L\beta\right)^i - \frac{1}{2\hat{\alpha}}\psi^{-10}\left(\hat{L}\beta\right)^{ij}\hat{\nabla}_j\ln\hat{\alpha} - \frac{2}{3}\psi^{-4}\hat{\nabla}^i K, \\
& -\frac{1}{2\hat{\alpha}}\psi^{-10}\hat{\alpha}\hat{\nabla}_j\left(\frac{1}{\hat{\alpha}}\hat{u}^{ij}\right).
\end{aligned} \tag{3.11}$$

Pfeiffer et al [102] have investigated Kerr-Schild-like binary initial data sets for equivalent physical situations, and have concluded that, at least for certain classes of data, solving the constraints via the CTS decomposition yields the least "excess" radiation on the initial data slice. However, the CTS decomposition is best suited to stationary or quasi-stationary systems (as with, for example, an equal-mass binary system in quasi-equilibrium). The head-on simulations we investigate here do not fit into this category, so we will content ourselves for the moment with the CTT decomposition.

3.3 Solving the Constraints in PETSc

The solver used comes from the PETSc libraries. It solves the four coupled non-linear constraint equations using a Newton method wrapped around a conjugate-gradient solver for the linear problem. If we write the constraints as a (non-linear) vector equation (using

[103], chapter 9 for notation):

$$\mathbf{F}(\mathbf{x}) \quad = \quad \mathbf{0},$$

then in a Taylor series sense, we can write "nearby" values of $\mathbf{F}$ as

$$\mathbf{F}(\mathbf{x}+\delta\mathbf{x}) \quad = \quad \mathbf{F}(\mathbf{x}) + \mathbf{J}\delta\mathbf{x} + O(\delta\mathbf{x}^2),$$

where $\mathbf{J}$ is the Jacobian matrix. Assuming that the "perturbed" value of $\mathbf{F}$ is zero, to linear order, the necessary $\delta\mathbf{x}$ that will achieve this will obey

$$\delta\mathbf{x} = -\mathbf{J}^{-1}\dot{\mathbf{F}}(\mathbf{x}). \tag{3.12}$$

This is the essence of an iteration of the Newton solver; a Newton line-search method, with cubic backtracking is used to calculate an appropriate $\delta\mathbf{x}$ from (3.12). Within this non-linear solve, we have a matrix inversion problem. This linear problem is handled by the *generalized minimum residual* method (see [103], Section 2.7 for a description), with a Block-Jacobi preconditioner. The Jacobian matrix is calculated numerically.

This is not a particularly efficient scheme, especially given the unigrid nature of the code. For a spatial domain containing N points altogether, the Jacobian matrix will be $4N \times 4N$ (as there are four constraint equations). For example, a cubical grid of side 100 means a Jacobian of side 4×10^6. The solver in our simulations has not been used successfully for domains larger than $(130)^3$ or so. For larger-domain runs or convergence runs in the same physical domain, a better solver method must be found.

The boundary conditions used for the solver are as follows: for the outer boundary, we use Robin boundary conditions, assuming a $1/r$ fall-off of all fields:

$$\begin{aligned} \psi &\rightarrow 1 + \frac{C^0}{r}, \\ V^i &\rightarrow 0 + \frac{C^i}{r} \end{aligned} \tag{3.13}$$

It should be pointed out that "proper" Robin conditions for a vector potential such as V^i are not as simple as those presented here [142]. However, the inconsistency of these boundary conditions was felt to be sufficiently unimportant to neglect.

For the inner (excision) boundary, we use Dirichlet conditions:

$$\begin{aligned} \psi &\rightarrow 1, \\ V^i &\rightarrow 0 \end{aligned} \tag{3.14}$$

This may be a good condition if we are dealing with data which is essentially single-hole data

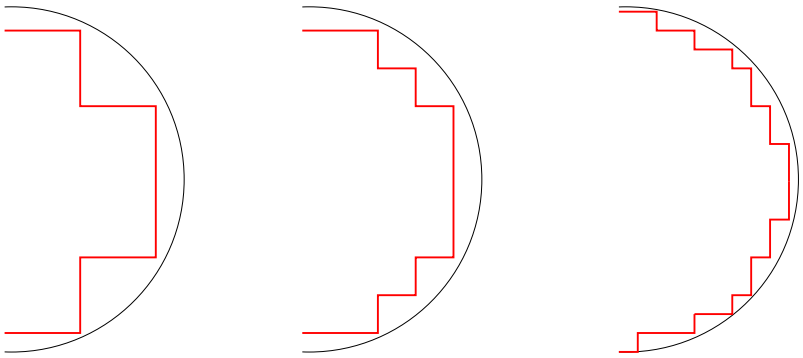

Figure 3.2: A "Lego sphere" excision mask at different resolutions, and how it fits within the ideal excision radius. Solver boundary conditions are applied at the grid points on the mask boundary, rather than at the excision radius.

near each inner boundary: the constraint violation will be very mild, and the Dirichlet data above essentially means that no correction is needed for the constraints on these boundaries. In a close-separation situation, the approximation will be less justifiable. Also, if we move the inner boundary (that is, the excision boundary) out or in, this affects the boundary conditions without affecting the actual "initial guess" data.

There is a related issue with using Dirichlet conditions on an excision boundary. Moving from a coarser resolution to a finer one, the excision boundary will change its shape, conforming better to the ideal position given by the excision radius (see Fig 3.2). This means that the same Dirichlet conditions will be applied to different physical positions. This should have an impact on the convergence of the solution potentials near the excision region.

A better idea for an inner boundary condition might be a Robin-like fall-off condition that enforces the Dirichlet values (3.14) only *at* the singularity. At present, such a condition has not been applied in the Maya code. A related idea is to use background data that is guaranteed to satisfy the constraints in a region containing the excision region. In that case, the solver boundary conditions could be taken as Dirichlet with impunity. Data of this sort is being investigated, but was not used for evolutions in the current work.

As an example of the results of solving initial data, I show below the post-solve potentials and Hamiltonian constraint for small-separation data using these boundary conditions, plotted along the z-axis. Here there is a single cubical excision region, with an ideal excision radius of $2.0m$.

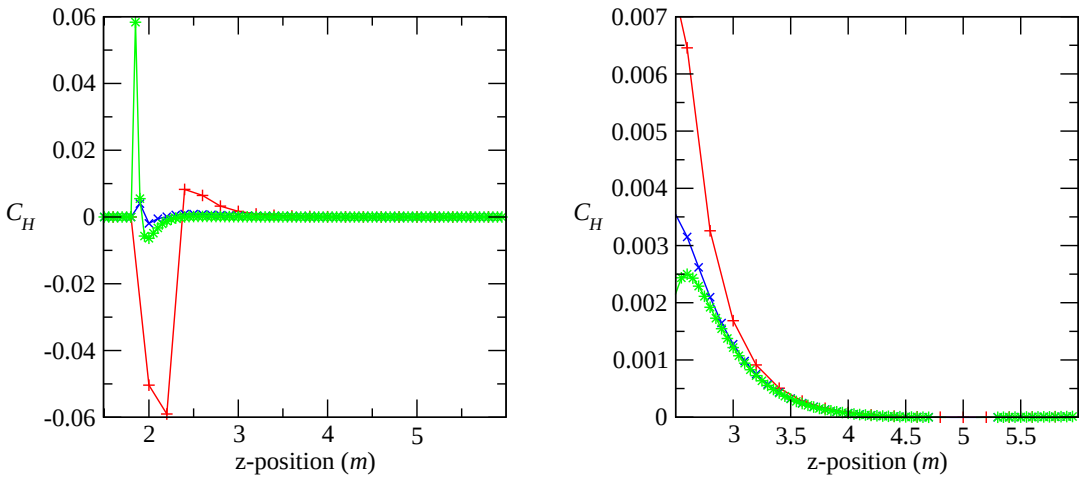

Figure 3.3: (i) Constraint violation outside solved small-separation data for different resolutions. Compare with Fig 3.1. (ii) The same violations away from the excision boundary. Medium and fine resolutions are scaled by factors of 4 and 16, respectively, to indicate the extent of convergence.

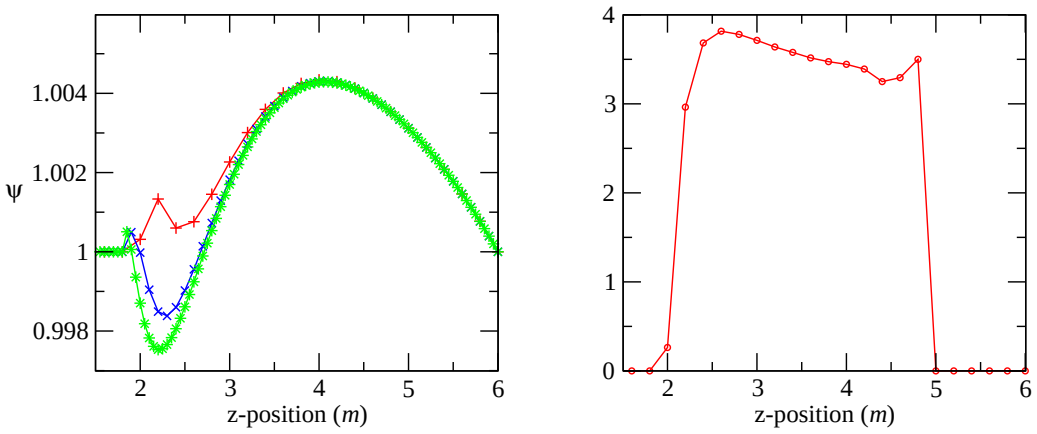

Figure 3.4: (i) Solution CTT factor ψ for different resolutions. (ii) Three-level convergence factor for ψ.

3.4 Binary Slicing

Solving the constraints (at least in the CTT and PTT cases) tells us nothing about the appropriate slicing conditions for this system. In lieu of a better idea, we construct something reasonable from the lapse and shifts of the individual Kerr-Schild holes. A single-hole solution has a lapse and shift in the form:

$$\alpha^{-2} = 1 + 2H(\ell_0)^2, \tag{3.15}$$

$$\beta_i = 2H\ell_i\ell_0, \tag{3.16}$$

where for zero spin, $H = \frac{M}{r}$, $\ell_\mu = \left(1, \frac{x}{r}, \frac{y}{r}, \frac{z}{r}\right)$. This form is invariant under boosts – one of the advantages of the Kerr-Schild coordinates. Specifically, under a boost v along the z-axis, all that changes is $H \to \frac{M}{R}$ and $\ell_\mu \to \left(\frac{\rho}{R}, \frac{x}{R}, \frac{y}{R}, \frac{\zeta}{R}\right)$, where $\rho \equiv \gamma(R - vZ)$ and $\zeta \equiv \gamma(Z - vR)$ (see Appendix, Section A.3). For small boosts, we can expand the lapse and shift in powers of v:

$$\alpha^{-2} = 1 + \frac{2m}{r} - \frac{4mz}{r^2}v + \frac{m(z^2 + 2r^2)}{r^3}v^2, \tag{3.17}$$

$$\beta_i = \left[\frac{2m}{r^2} - \frac{2mz}{r^3}v + \frac{m(r^2 - 2z^2)}{r^4}v^2\right]\left(x, y, z - vr + v^2 z\right). \tag{3.18}$$

Different possible binary slicings can be tried, constructed in terms of the Kerr-Schild form. The simplest conditions that respect the single-hole limit are:

$$\alpha = {}^A\alpha\,{}^B\alpha, \tag{3.19}$$

$$\beta^i = {}^A\beta^i + {}^B\beta^i. \tag{3.20}$$

These conditions, however, do not have the "close limit" – where the separation of the holes is very small, they do not look like the lapse and shift of a single IEF hole. For instance, for equal masses, and equal but opposite small boosts, the "total" lapse will be

$$\alpha = \frac{r}{r + 2m} - \frac{m(z^2 r - 6mz^2 + 2r^3 + 4mr^2)}{r(r + 2m)^3}v^2 + O(v^4). \tag{3.21}$$

A compound slicing that has better limits is given by:

$$\alpha = \left({}^A\alpha^{-2} + {}^B\alpha^{-2} - 1\right)^{-\frac{1}{2}} \tag{3.22}$$

$$\beta^i = \gamma^{ij}\left({}^A\beta_j + {}^B\beta_j\right) \tag{3.23}$$

In the close limit, when two holes with equal masses, and equal (but opposite) boosts,

coincide at the origin, the compound lapse is given by:

$$\alpha = \sqrt{\frac{r}{r+4m}}\left[1 - v^2\frac{2m}{r+4m}\left(1 - \frac{z^2}{2r^2}\right)O(v^4)\right]. \tag{3.24}$$

Here we see that the dominant behavior is exactly what we would expect for a single IEF hole of mass $M = 2m$, centered at the origin. However, at the lowest boost order, $O(v^2)$, there is a correction to the total mass (that is, a "monopole" term with no directional dependence), and also considerable anisotropy.

Of course, as these fields are pure gauge, the exact form is not critical to the physics of the problem. A correct close limit is desirable, but hardly necessary. What *is* necessary is that for initially separated holes, the lapse and shift should locally resemble the Kerr-Schild lapse and shift of a single boosted hole. Both of these choices satisfy this requirement. For computational efficiency during evolution, we generally pick the former, (3.19)–(3.20).

3.5 Diagnostics: Apparent Horizons

To complete this chapter, I would like to mention briefly some of the general diagnostic methods used in analyzing evolutions in Chapters 4 and 5.

The best diagnostic of the state of a black hole is the event horizon. This gives us the true causal boundaries of the space-time as a whole. Event horizon finders exist for numerical codes; most notably, Diener [56] has recently made one available for use with the Cactus computational toolkit. However, this and similar finders are global in nature (as are event horizons themselves), and must be run backward in time after the end of the evolution. Additionally, they entail considerable data storage, and are not practical at present for the Maya code. For reasons of practicality, therefore, we have not yet attempted event horizon extraction.

Instead, we look at the poor man's event horizon, the *apparent horizon*. This is defined as the outermost *marginally trapped surface*; that is, the outermost surface where the expansion of outgoing null vectors vanishes (see [81]). To be more specific, if we take a 2-sphere S in our spatial hypersurface, the apparent horizon condition can be expressed as an elliptic equation, the vanishing of the expansion κ on this surface:

$$\kappa \equiv D_p n^p + K_{pq} n^p n^q - K = 0. \tag{3.25}$$

Here, n^i is the outward-pointing spatial normal to the two-surface S, and all quantities are intrinsic to the hypersurface. In this sense, the apparent horizon is a special case of the more general surface of constant expansion.

For a stationary single-BH space-time, the intersection of the event horizon with any spatial hypersurface (that is, "at any time") will coincide with an apparent horizon on that

hypersurface. In the more general situation where radiation is present, an apparent horizon will lie strictly *inside* the event horizon.

The apparent horizon is, unfortunately, not a coordinate-invariant concept. Even for single, stationary BHs, it is possible to slice the space-time such that an apparent horizon cannot be found [101], although there will certainly be an event horizon. For this reason, there is general wariness about putting too much weight on the apparent horizons and associated quantities in serious numerical studies. Nevertheless, some of these quantities are suggestive, and supply guidance in the interpretation of the spatial slice. The main diagnostic we use is the presence of an apparent horizon to begin with. This tells us that there is a black hole in that region of the numerical grid.

After locating an apparent horizon, we look at some associated quantities:

- The *effective mass* of an apparent horizon is the mass implied by the horizon area formula. That is,

$$M_{AH} \equiv \sqrt{\frac{A_{AH}}{16\pi}}; \tag{3.26}$$

- The *horizon distortion* is defined here as the difference from unity of the ratio of proper equatorial circumference to proper polar circumference. That is,

$$\epsilon_{AH} \equiv \frac{C_E}{C_P} - 1 \tag{3.27}$$

- Finally, a measure of the separation of two black holes on a hypersurface is given by the shortest proper distance between the two apparent horizons. For the simulations presented here, axisymmetry guarantees that this will be along the z-axis. For HuMaSh data – ignoring the issue of the constraints, and assuming the holes are far enough apart that the individual apparent horizons are a radial distance $2m$ from the singularities – the horizon-horizon proper separation D will be (see Appendix Section B.4):

$$D = 2\int_{z=0}^{d/2-2m} \sqrt{1 + \frac{2m\gamma(1-v)^2 d}{d^2/4 - z^2}}\, dz, \tag{3.28}$$

  where the holes are boosted toward each other with velocity $\pm v$, and corresponding Lorentz factor γ. The integral has to be calculated numerically, and for large separations, will converge to $d - 4m$.

The main use of the apparent horizon effective mass is in binding energy calculations; the difference between the total mass of the system (represented by the ADM energy, for instance – see section B.1), and the intrinsic mass of the holes (the sum of the effective

masses here) gives a mass deficit, or binding energy analogous to the binding energy in nuclear physics. See [27] for general use in initial data analysis, and [47] for use of a binding energy argument in obtaining the "innermost stable circular orbit" (ISCO).

There are two or three widely used algorithms employed to find the apparent horizon(s) in a spatial slice. The most robust of these, the *flow* methods, do not need a good initial guess for the horizon position, and can handle changes in topology (for instance, going from two single-hole apparent horizons to a deformed common horizon). The reader is directed to [82, 117, 131] for details of these and other methods.

For the most part, I used the CACTUS code's AHFinder module for horizon finding in the runs mentioned below; more specifically, I used the CactusEinstein implementation of curvature flow method. Although slow, the code has the virtue of stability – naturally it will fail if it encounters excised grid points, but it exits gracefully in these circumstances.

One area where we would like some measure of accuracy is in the determination of the apparent horizon area, and hence of the effective mass M_{AH}. The two controllable parameters in the CACTUS implementation of the flow method are the angular resolution of the mesh and ℓ_{max}, the maximum angular mode included in a spherical harmonic decomposition of the horizon. For a single IEF hole, the effective mass obtained from the AHFinder is accurate to within 0.2% for the default parameters (i.e., a 100×100 point mesh, and $\ell_{max} = 8$). The accuracy stays high even if we increase the excision radius, up until $R_h = 1.9m$, when the finder procedure fails at the resolution used ($\Delta = 0.2m$).

Unfortunately, it is difficult to think of a more rigorous test case for the finder. It is true that an apparent horizon is a local entity; this makes it very susceptible to local errors and noise.

The apparent horizon is hardly the last word in local black hole analysis techniques. Recently, much classical work has been done on more sophisticated related ideas. Most notably, Ashtekar et al have introduced the concept of the *isolated horizon*, an $S^2 \times R$ null surface which is not crossed by matter or radiation over a finite time period. This condition can be satisfied if the surface is both expansion- and shear-free. Thus an isolated horizon will always be an apparent horizon (or rather, its intersection with a particular time slice will be), but the converse does not hold. Being able to find an isolated horizon allows us to measure the black hole mass and spin much more accurately than before. The isolated horizon framework has been implemented in the MAYA code [57]; though not used in this thesis, it may become important for grazing collision simulations in the future.

Another, more recent development is Ashtekar and Krishnan's *dynamical horizons* [13]. These are space-like surfaces that limit to isolated horizons when no radiation is crossing them. As these are more general than isolated horizons, and we can be sure that a numerically generated space-time will be awash with radiation until very late times, these objects are of great interest. This framework is new and has not seen much numerical

implementation to date; I know of only one recent work that has used it [14].

3.6 Diagnostics: ADM Energy

Apparent horizons are local geometric entities, identifying and describing the space-time near individual black holes.

An important measure of the *global* energy content of a spatial hypersurface is the ADM energy. Assuming that the space-time is asymptotically flat, it can be calculated as a surface integral at infinity. The "traditional" formula can be found in Chapter 11 of Wald [135]; it, and the associated global ADM three-momentum, are defined in his equations (11.2.14)–(11.2.15) as surface integrals taken over a sphere at spatial infinity:

$$\begin{aligned} E_{ADM} &\equiv \lim_{R\to\infty} I_{trad} \\ &= \frac{1}{16\pi} \oint_{r=R} \left(\gamma_{mn,m} - \gamma_{mm,n}\right) N^n dA, \\ p_n &= \lim_{R\to\infty} \frac{1}{8\pi} \oint_{r=R} \left(K_{mn} N^m - K^m_{\ m} N_n\right) dA. \end{aligned}$$

Implicit in these formulae is the expectation that the space-time is asymptotically flat; that is, all fields drop to flat space at spatial infinity sufficiently quickly. $N^j \equiv \frac{x^j}{r}$ here is just the flat-space normal to the surface, and $dA \equiv r^2 d\Omega$ is the euclidean area element.

I have implemented the standard ADM formula in the Maya code, as a summation over a surface of points, equally spaced in θ and ϕ. Of course, with a finite Cartesian mesh, we cannot take the limit $R \to \infty$, so instead we choose a smooth surface as close to the outer boundary as possible. Additionally, this surface will not always be spherical, if the computational domain admits an ellipsoidal surface significantly larger than the maximum possible spherical boundary.

3.6.1 Systematic Errors in ADM Energy Calculation

It is interesting to note that using the exact IEF solution in the standard ADM integral yields a surprisingly simple answer:

$$\begin{aligned}
\gamma_{ij} &= \delta_{ij} + \frac{2Mx_ix_j}{r^3} \\
\Rightarrow \gamma_{ij,k} &= \frac{2M}{r^3}\left[\delta_{ik}x_j + \delta_{jk}x_i - \frac{3x_ix_jx_k}{r^2}\right] \\
\Rightarrow \gamma_{ij,i} &= \frac{2Mx_j}{r^3}, \quad \gamma_{ii,j} = -\frac{2Mx_j}{r^3} \\
\Rightarrow I_{trad} &= \frac{1}{16\pi}\oint_{r=R}\frac{4Mx_jx^j}{r^4}r^2d\Omega \\
&= \frac{1}{16\pi}4M4\pi = M.
\end{aligned}$$

That is, the ADM surface integral yields the exact BH mass, no matter what the radius of the coordinate sphere. It should not be necessary to produce a sphere at infinity to get the exact ADM energy, assuming high enough resolution. Thus, inaccuracies observed should be due entirely to the finite-differencing of spatial derivatives, and the interpolation of data onto the surface, both second-order-accurate operations in Maya.

This is a special case; other coordinatizations of the Schwarzschild solution – for example, the exact Isotropic Schwarzschild solution – have no special relationship with the ADM integral:

$$I_{trad} = M\left(1 + \frac{M}{2R}\right)^3.$$

Thus, taking a finite-radius surface integral will result in a large systematic error, unconnected to the resolution of the surface itself.

To test these observations, I obtained the ADM energy for both IEF and Isotropic Schwarzschild initial data sets, with different angular resolution on the two-sphere surface, and at different radial distances. The results can be found in Section B.1; they largely confirm the expectations above. In particular, the radial distance had no measurable effect on the IEF integral, assuming a fixed angular resolution. For a 48×96 angular grid, the resolution error was a little under 3%.

In contrast, the result of the ADM integral showed a very strong radial dependence for Isotropic Schwarzschild – almost 19% for a radius of $10M$. Moreover, after adjusting for the constant resolution error, the remainder conformed well to the expectation of position dependence.

3.6.2 Corrections to Fall-off Error

The examples above show that, in most cases, one cannot expect the full ADM energy to result from a finite-radius surface integral. However, one can improve these numbers immensely; by calculating the integral at two different radial distances, one can cancel the

inevitable $1/R$ error.

Assume that the finite-R integral represents the actual ADM energy plus some error that falls off linearly with R: $I(R) = E + BR^{-1}$. Then we can take two different R values and eliminate B:

$$\begin{aligned} I_1 &= E + BR_1^{-1} \\ I_2 &= E + BR_2^{-1} \\ \Rightarrow B &= (I_1 - I_2)\frac{R_1 R_2}{R_2 - R_1} \\ \Rightarrow E &= \frac{I_2 R_2 - I_1 R_1}{R_2 - R_1} \\ &= \frac{I_2 - \kappa I_1}{1 - \kappa}, \end{aligned}$$

where $R_1 = \kappa R_2 < R_2$. This cancellation is carried out in the MAYA code, with a default value $\kappa = 0.75$.

The result of this adjustment, in the far-field regime, is to cut the ADM error hugely. However, in the near-field (as we are), it is not clear whether this improves matters or compounds the error. Given our general ignorance of the real character of a numerically evolved space-time, we must take the angular resolution error (3% for 48×96) as a minimum error inherent in the determination, and seriously consider allowing for much greater error, such as the 19% mentioned above, if the form of the three-metric is not exactly that of Kerr-Schild.

3.7 Diagnostics: Weyl Scalars

The five complex Weyl scalars Ψ_0, Ψ_1, Ψ_2, Ψ_3, and Ψ_4 contain the ten independent degrees of freedom in the Weyl curvature tensor. They are calculated in the Newman-Penrose formalism, using a null tetrad adapted to the symmetry of the underlying space-time. In the Newman-Penrose formalism, we first choose a tetrad consisting of two real ($\mathbf{l}$ and $\mathbf{n}$) and one complex ($\mathbf{m}$) null vectors with certain additional orthogonality conditions:

$$\mathbf{l} \cdot \mathbf{n} = -1 \ , \ \mathbf{m} \cdot \bar{\mathbf{m}} = 1, \tag{3.29}$$

with all other products zero.

We then cast all tensors in tetrad form by forming products with combinations of the tetrad vector. In particular, the Weyl tensor's ten degrees of freedom can be neatly expressed

as five complex scalar fields, written as $\Psi_0 \rightarrow \Psi_4$:

$$\begin{aligned}
\Psi_0 &= C_{\mu\nu\rho\sigma} l^\mu m^\nu l^\rho m^\sigma \\
\Psi_1 &= C_{\mu\nu\rho\sigma} l^\mu n^\nu l^\rho m^\sigma \\
\Psi_2 &= C_{\mu\nu\rho\sigma} l^\mu m^\nu \bar{m}^\rho n^\sigma \\
\Psi_3 &= C_{\mu\nu\rho\sigma} l^\mu n^\nu \bar{m}^\rho n^\sigma \\
\Psi_4 &= C_{\mu\nu\rho\sigma} n^\mu \bar{m}^\nu n^\rho \bar{m}^\sigma .
\end{aligned} \tag{3.30}$$

I have written a code in MAYA to calculate these, and related, quantities using a method outlined by Gunnarsen et al [76], employing the ADM "3+1" quantities, γ_{ij} and K_{ij}. The procedure is outlined in the Appendix, Section B.2. I shall comment here only on the uses of the resulting scalars.

The real usefulness of the Weyl scalars comes in the analysis of the fundamental properties of a space-time; the Petrov classification of the space-time can be investigated via the properties of these scalars under transformations. For instance, if we can find a null tetrad where the only non-vanishing Weyl scalar is Ψ_2, then the space-time is of Petrov type D (see Chap 4 of [89]).

Once we have calculated the Weyl scalars, we can deduce the value of the invariants I (also known as the Kretschmann invariant), J, and the "speciality index" S:

$$\begin{aligned}
I &\equiv \Psi_0\Psi_4 - 4\Psi_1\Psi_3 + 3\Psi_2^2, \\
J &\equiv \begin{vmatrix} \Psi_4 & \Psi_3 & \Psi_2 \\ \Psi_3 & \Psi_2 & \Psi_1 \\ \Psi_2 & \Psi_1 & \Psi_0 \end{vmatrix}, \\
S &\equiv \frac{I^3}{27J^2}.
\end{aligned} \tag{3.31}$$

For the IEF or Schwarzschild forms of the Schwarzschild solution, all the Weyl scalars are identically zero, except for

$$\mathrm{Re}(\Psi_2) = \frac{M}{r^3},$$

while the invariants defined above will be

$$I = \frac{3M^2}{r^6} \quad , \quad J = -\frac{M^3}{r^9} \quad , \quad S = 1.$$

In fact, the invariant S will be unity if the space-time is Petrov Type D, whatever coordinates we're using – it is an entirely coordinate-free measure of the stationarity of a space-time. Unfortunately, as S has only one distinguished value that I am aware of – unity – and is dimensionless, there is no obvious measure of what constitutes a large deviation from type

D. Baker & Campanelli [15] treat S as representing significant perturbation from a type-D state if it differs from unity by a factor of two either way.

Apart from fundamental properties of the full space-time, the Weyl scalars are useful in perturbative treatments. In the early 1970s, Teukolsky & Press [129] worked in the Newman-Penrose formalism to derive equations describing the perturbation of a Kerr black hole by scalar, electromagnetic, and gravitational waves. In the latter case, the perturbation quantities were the Weyl scalars, Ψ_0 and Ψ_4. The fall-off properties of the solutions led to the characterization of Ψ_0 as representing radiation falling into the black hole, and Ψ_4 as outgoing radiation. These calculations are gone through in their full gory detail by Chandrasekhar [42].

Finally, I should mention that Beetle & Burko [20] have suggested a new space-time scalar, calculated from on-time-slice quantities, that could be useful for the analysis of the gravitational-wave content of the space. This uses S in its construction, but allegedly has better understood properties in perturbative situations. This scalar has not been implemented in Maya to date.

3.8 Diagnostics: Quasi-normal Mode Analysis

Although the merger regime is not well understood, we know that the final stage of a merger must be ring-down, when the underlying black hole shakes off the last of the distortions around it. The ways in which it can do this are limited – a Schwarzschild black hole has a set of well-defined quasi-normal modes (QNM), "quasi" because they have a complex frequency, with a guaranteed damping rate as well as observable physical frequency.

More will be said about QNM for Kerr black holes in Chapter 6, but for now, we note that the modes have two "quantum numbers" – n and ℓ – for a Schwarzschild BH[2]. For each ℓ, the modes are ordered so that the fundamental mode, $n = 1$, is the most slowly damped.

Table 3.1 gives some of the most fundamental QNM frequencies for a Schwarzschild BH, together with the implied physical period for a BH of mass M[3]. These are the frequencies that we expect a ringing black hole to exhibit. However, we don't know how much energy will go into each mode; our experience leads us to be sure only that the $n = 1, \ell = 2$ mode will contain the bulk of the radiated power. We *do* know, though, that in simulations containing the symmetries used here, only even-ℓ modes can be excited.

Being able to identify the longest QNM period (or shortest finite frequency) in a ring-down spectrum will give us another way to determine the underlying mass of the final object. An attractive aspect of this method of mass determination is its global nature – the QNM determination uses information taken over several hundred m to identify the fundamental

[2]For a Kerr black hole, these modes are split further, in analogy with atomic spectra.

[3]I am indebted to K. Kokkotas for supplying many of the higher-ℓ modes, unavailable in current literature.

(n, ℓ)	ω	Period (M)	Damping time (M)
(1,2)	$0.3737 - i0.0890$	16.81	11.24
(2,2)	$0.3467 - i0.2739$	18.12	3.65
(3,2)	$0.3011 - i0.4682$	20.87	2.14
(1,3)	$0.5994 - i0.0927$	10.48	10.79
(2,3)	$0.5826 - i0.2813$	10.78	3.55
(3,3)	$0.5517 - i0.4791$	11.39	2.09
(1,4)	$0.8092 - i0.0942$	7.76	10.62
(2,4)	$0.7966 - i0.2843$	7.89	3.52
(3,4)	$0.7727 - i0.4799$	8.13	2.08
(1,5)	$1.0123 - i0.0949$	6.21	10.54
(1,6)	$1.2120 - i0.0953$	5.18	10.49
(2,6)	$1.2036 - i0.2866$	5.22	3.49
(3,6)	$1.1871 - i0.4806$	5.29	2.08
(1,7)	$1.4097 - i0.0955$	4.46	10.47
(1,8)	$1.6062 - i0.0957$	3.91	10.45
(1,9)	$1.8018 - i0.0958$	3.49	10.44
(1,10)	$1.9968 - i0.0959$	3.15	10.43
(1,11)	$2.1913 - i0.0959$	2.87	10.43
(1,12)	$2.3855 - i0.0960$	2.63	10.42

Table 3.1: Selected quasi-normal mode frequencies for Schwarzschild black holes.

frequency. If we are certain in our identification of this mode, the result should be more robust than more local determinations such as the ADM mass (a surface integral near the outer boundary) or the apparent horizon effective mass (a surface integral near the horizon).

If we can identify other modes, we will have much more knowledge about the manner in which head-on mergers excite the final black hole. Notice from Table 3.1 that the fundamental mode for each ℓ is *not* the one with the largest period. It is possible that the largest period visible will be that of, say, the first overtone of $\ell = 2$; we usually rely on the much greater damping (as evidenced by the larger imaginary component of the frequency) to reduce its impact in any observed spectrum.

In practice, we will identify the frequencies present by taking a Fourier transform (FT) of the observed wave form[4]. For an undamped signal, the ideal FT would be a delta function – infinitely high and infinitely narrow. Of course, we will not expect to see this in an actual spectrum. For quasi-normal modes, an important source of line broadening will be the intrinsic damping of the mode itself – see the Appendix Section B.3 for details. For a single QNM mode with amplitude A, frequency f and damping time τ, the real part of the Fourier

[4] In principle, given a perturbed field h, the correct way to identify modes with different ℓ is to integrate h against spherical harmonics. For simplicity here, however, I restrict myself to looking at a suitable field scalar at a single detector position only.

transform of the signal will be

$$Re\left[\tilde{s}(\nu)\right] = \frac{A\tau}{1 + 4\pi^2(f-\nu)^2\tau^2}. \tag{3.32}$$

From this shape we can extract redundant information about the black hole mass. The peak position of the transform will be

$$\nu^* = f,$$

while the breadth of the can be measured by a "half-peak" width at

$$\nu_{\frac{1}{2}} = f \pm \frac{1}{2\pi\tau}.$$

One can construct similar formulas for a "two-thirds-width", a "three-quarters-width", and so on. These might be necessary if a peak's lower reaches are buried in a background or the tails of neighboring peaks.

Thus, by measuring the peak position and the width, we can extract the frequency and damping time, and assuming the correct identification of the mode, we can deduce the background mass from Table 3.1.

3.9 Diagnostics: A Final Word

I have looked here at several different methods of extracting physical information about the evolution system. None of these is perfect; their diversity is their main strength. The different measures work with different aspects of the evolution fields. Some are very local – M_{AH} and M_{ADM} – while others are "quasi-global", such as the use of the Ψ_2 profile, or the QNM fundamental frequency fit. When trying to determine the same information from diverse sources, we at least have several different methods of attack. The convergence of the results would be a strong heuristic test of the accuracy of our evolutions.

Having addressed the initial data sets to be employed, and the diagnostics available to me, in the next chapter, I will move on to actual evolutions.

Chapter 4

Small Initial Separation: Distorted Black Holes

Preparatory to attempting the full head-on merger of two black holes, we might first use similar data to tackle a conceptually simpler problem – the full 3D evolution of small-separation data. This is data where the initial coordinate separation of the singularities is sufficiently small that a single common apparent horizon is present at the initial time slice, and so we are justified in using a single excision region and treating the space-time as that of a distorted single black hole. There are at least two good reasons for pausing here before going on to the full merger:

- First, looking at the small-separation problem before the full merger allows us to distinguish between numerical issues that are generic in highly non-linear numerical problems (such as live slicing conditions and boundary conditions) and the additional problems involved in dealing with a merger simulation (dynamic excision, merger procedure, analytic slicing).

- Second, distorted black holes are interesting in themselves. close-limit evolutions have certainly been performed in axisymmetry [9, 10], where the Misner data used could be compared to perturbative calculations by Price & Pullin [106] and others. In full 3D, however, most distorted simulations have been approached as a study of direct perturbations: a single black hole overlaid with wave packets of some kind; Brill waves, for instance as in [39]. Although very restrictive in its symmetry, binary distorted data should be of a very different – and possibly more violent – nature to these wave overlays.

This chapter is laid out as follows. The "distortion" assumption can be made when the singularities are not particularly close, and considerable distortion exists in the initial data. In section 4.1, I discuss the critical point dividing the binary and single-BH cases on the initial time slice.

To carry out an evolution, we need the initial data of Chapter 3; we must also decide on slicing and outer boundary conditions. These will be addressed in Section 4.2.

Given these choices, I present small-separation evolutions in Section 4.3, and attempt to analyze the physical content of the results.

Finally, in Section 4.5 I discuss what has been learned from these runs, and what light they throw on the full merger runs of Chapter 5.

4.1 The Point of BBH Fusion

The first step in carrying out a simulation of the merger of two well-defined black holes is obtaining initial data that represents two well-defined black holes. Of course, the critical coordinate separation at which this distinction arises depends strongly on the form of the initial data. All other things being equal, we may assume that the larger the initial separation, the greater the distortion of the final black hole. The only real way to determine if one or two black holes are present on a particular time slice is to evolve for the entire history of the space-time, and then trace the final event horizon backwards in time to the slice we are interested in. Nevertheless, these procedures would necessitate carrying out full test evolutions to late times, in order to find a non-disjoint event horizon that would justify our performing the evolutions in the first place.

Lacking an event horizon, we analyze the initial data through the presence of apparent horizons (see Section 3.5). We investigate a sequence of initial-data slices, each more separated than the last, until our apparent horizon finder fails to find a common apparent horizon. The results of this procedure are presented in Table 4.1. The results are for unboosted and strongly boosted ($v = 0.5$) holes. In the latter case, the presence of boosts increases the energy contained in the space, and will arguably lead to a larger and more spherical apparent horizon. Where an apparent horizon is found, the ratio of the proper length of the equator and polar meridian is given, as a measure of the coordinate-independent distortion of the apparent horizon.

The solver finds a common apparent horizon for $d = 5.75m$ only for boosted data, and no common apparent horizon for any boost for $d \geq 6.0$. This is not a guarantee that no such apparent horizon exists; there could be one there, highly distorted – perhaps even locally concave, like a peanut. In addition, there will in general still be a single-BH event horizon present. Thus to be sure that we really are merging two well-defined BH, we must ensure that the coordinate separation is well over $6m$.

4.2 Slicing and Boundary Conditions

For long-term stable evolutions, we know that our slicing must be dynamic. The minimal way to achieve this is through the use of an analytic densitized lapse (2.23): the physical lapse α is automatically dynamic as it is the product of an analytic Q and an evolved conformal factor $e^{-6n\phi}$.

However, in these evolutions, we consider this to be insufficient: we evolve the densitized

	$v = 0.0$				$v = 0.5$			
d_s	M_{AH}	C_E	C_P	$\epsilon(\%)$	M_{AH}	C_E	C_P	$\epsilon(\%)$
0.0	2.63	33.08	33.06	0.0	3.14	40.12	38.86	3.2
0.5	2.66	33.35	33.37	0.0	3.01	38.68	37.24	3.9
1.0	2.69	33.70	33.84	-0.4	2.84	36.61	35.26	3.8
1.5	2.69	33.51	33.90	-1.2	2.71	34.79	33.75	3.1
2.0	2.59	32.04	32.84	-2.4	2.64	33.63	32.96	2.0
2.5	2.51	30.63	32.01	-4.3	2.58	32.56	32.36	0.6
3.0	2.46	29.56	31.67	-6.7	2.53	31.57	31.99	-1.3
3.5	2.56	30.26	33.24	-9.0	2.92	36.07	36.96	-2.4
4.0	2.32	26.31	30.60	-14.0	2.45	29.44	31.46	-6.4
4.5	2.29	25.00	30.69	-18.5	2.53	29.68	32.79	-9.5
5.0	2.23	23.15	30.59	-24.3	2.38	26.83	31.22	-14.1
5.5	2.20	21.38	31.08	-31.2	2.33	25.20	31.16	-19.1
5.75	—	—	—	—	2.31	24.36	31.25	-22.0
6.0	—	—	—	—	—	—	—	—

Table 4.1: Survey of common apparent horizons surrounding holes of varying coordinate separations d_s. All dimensionful quantities are expressed in units of the single-hole "bare mass", m. Horizons are found using the flow algorithm. v is the boost velocity, M_{AH} the effective apparent horizon mass, C_E the proper equatorial circumference, C_P the proper polar circumference, and $\epsilon \equiv C_E/C_P - 1$ is the apparent horizon distortion. "—" indicates the finder could not locate an apparent horizon.

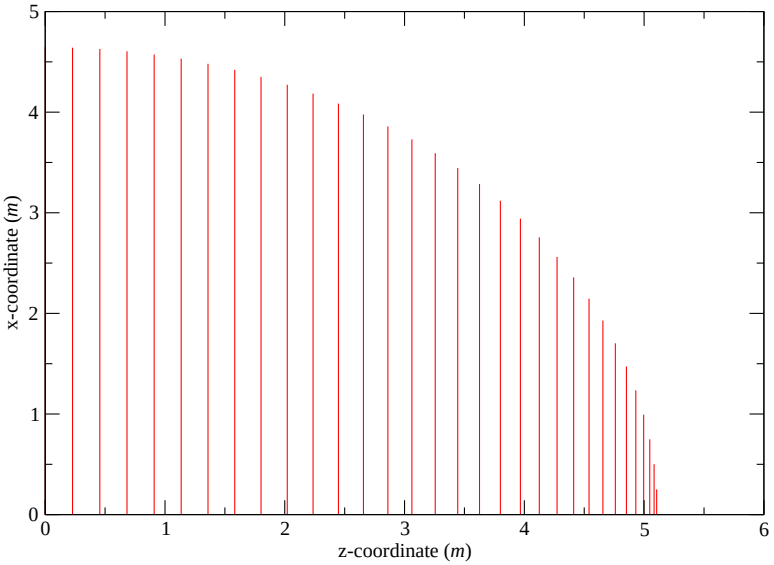

Figure 4.1: Overhead view of apparent horizon for on-top data with initial singularity separation of $d_s = 3.0m$.

lapse and shift using the 1+log (2.27) and Gamma-driver (2.28) conditions. The reason for this has to do with coordinates. It may be true that an evolution with analytic densitized lapse will be stable on its own; and no doubt the end result will be a steady-state IEF (or Kerr) black hole. But if we keep any field in the binary data form – even gauge fields such as the densitized lapse and the shift – the coordinates as a whole cannot settle down to the Kerr-Schild form; we will have frozen in an anisotropy, even if it is purely a coordinate one. For this reason, we use the live slicing conditions in our evolutions. This will give the slicing a chance to relax to a Kerr-Schild form for all fields (though of course, it does not guarantee that it will).

4.3 Small-Separation Evolutions

For our test, we take initially unboosted data with a center-to-center coordinate separation of $d_s = 3.0m$. These evolutions were carried out on a domain of size $(15m)^3$, with a grid spacing of $\Delta = 0.2m$.

For this initial coordinate separation, we find a common apparent horizon with a distortion $\epsilon = -6.7\%$ (see Table 4.1). Fig 4.1 shows an "overhead view" of one octant of the common AH present on the initial slice. This is clearly distorted in coordinate space, though not necessarily to the same extent as the invariant ϵ.

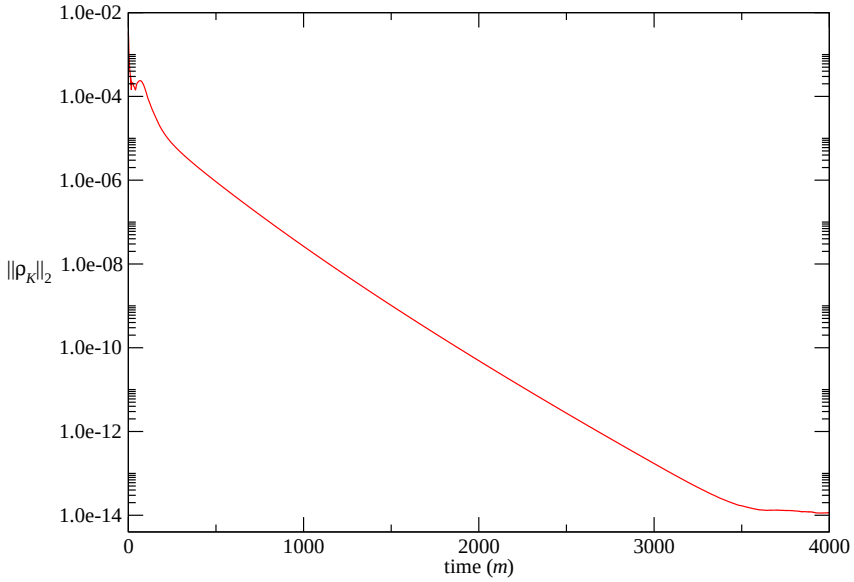

Figure 4.2: L2-norm of ρ_K, the right-hand-side of $\partial_t K$, for on-top simulation.

4.3.1 Stability

Our first requirement with these runs is that they should be long-term stable: we monitor the L2-norm of $\partial_t K$, to see if it drops to machine precision in a finite time. As long as our errors are above machine precision, there is always the chance that they are masking the development of a growing mode that might kill the simulation later. Machine precision itself is essentially the "quantum" of mode excitation in numerical terms – no mode can be excited with a smaller initial amplitude than this, so if we achieve this level, we can be fairly confident there are no unstable modes present[1].

That we do achieve long-term stability here is indicated by Fig 4.2, a semi-log plot of the L2-norm over time of ρ_K, the right-hand-side of $\partial_t K$ from (2.21). The exponential damping time for this system is $T_e \approx 200m$, until saturation at a level of about 10^{-14}. Machine precision itself is another two or three orders of magnitude lower than this level; the discrepancy is most likely to do with the manner in which ρ_K is calculated: each field value used to calculate each term in (2.21) possesses some error, and these will aggregate. That is, for ρ_K, a level of 10^{-14} is effectively machine precision.

[1] Of course, the runs presented here use only a restricted set of domain sizes, resolutions, etc. It is possible that we have been fortunate enough to exclude possibly unstable modes by accident.

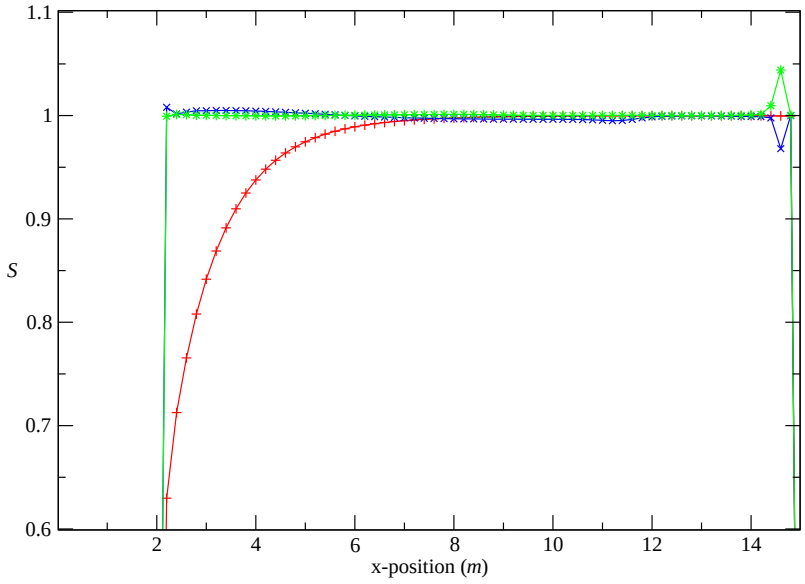

Figure 4.3: The Speciality Index S along the x-axis at times $0m$ (+), $10m$ ($\times$), and $20m$ ($*$).

Another measure of the stability of the system is the settling down of the "speciality index" S (3.31) at late times. Fig 4.3 shows S along the x-axis at times $t = 0m$, $10m$, and $20m$. Ideally, $S = 1$ for any type-D space-time (which includes all stationary black-hole solutions); as we can see, S quickly converges to unity, and is essentially frozen after a few hundred m of evolution.

4.3.2 Mass Determination

Now that I have established stability in these evolutions, I would like to extract physical information about the system. The simplest information to extract should be the bulk parameters describing the end-state BH, that is mass M and spin a (though of course, the symmetric head-on collision we have here involves no spin).

Fig 4.4 shows the ADM energy M_{ADM} (see Section 3.6) and apparent horizon effective mass M_{AH} (Section 3.5) for stable runs with outer boundaries at $15m$ and $30m$. After violent oscillation in the first $100m$ of evolution or so, both settle down to a steady state. Interestingly, M_{ADM} asymptotes to a much lower value – $1.33m$, as compared to $1.85m$ for M_{AH}. In a steady state, when all radiation initially present has either fallen into the hole

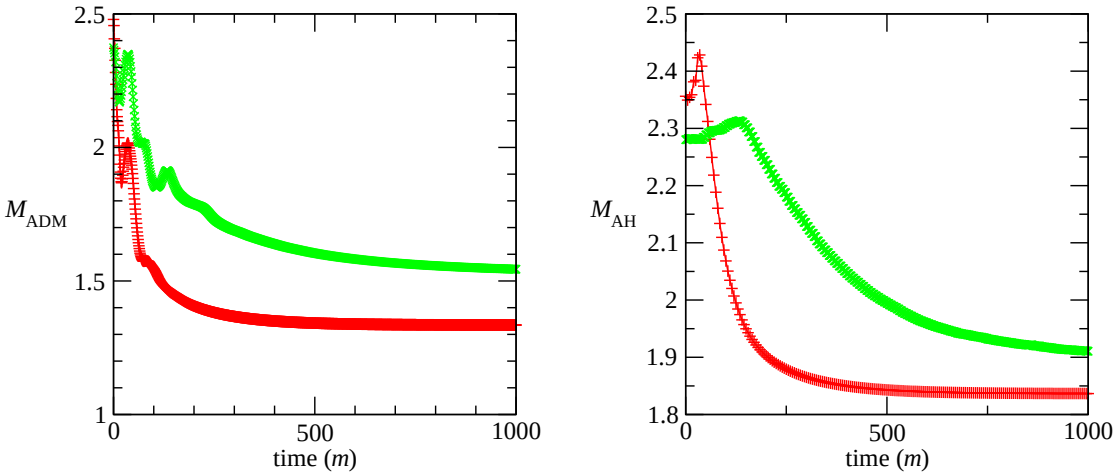

Figure 4.4: Apparent horizon effective mass and ADM energy for on-top data during evolution, for outer boundaries at $15m$ (red, +) and $30m$ (green $\times$).

or left the system, we would expect the two numbers to converge. Given the difference, we tend to trust the apparent horizon numbers more, if only because we are grossly violating the "coordinate sphere at infinity" requirement of the ADM condition.

Also shown in Fig 4.4 are the same quantities for an identical run in a larger domain ($(30m)^3$), but with slightly coarser resolution ($\Delta = 0.25m$). The qualitative result is the same: M_{ADM} and M_{AH} settle down to non-converging steady-state numbers. The end-state $M_{ADM} = 1.53m$, while $M_{AH} = 1.89m$. The greater uniformity of the apparent horizon numbers lends them greater credence.

An alternative measure of the mass of the final hole depends on the coordinates of the final solution. If the coordinates are close to Kerr-Schild post-merger, we know that the only non-zero Weyl scalar at late times should be the monopole Ψ_2, and that for a BH of total mass M, $\Psi_2 = \frac{M}{r^3}$. Fig 4.5 shows the late-time form of Ψ_2 along the x- and z-axes. Also shown is a fit to the functional form $\frac{M}{r^3}$. Fitting to the x-axis profile for $x \in [3.5m, 14m]$

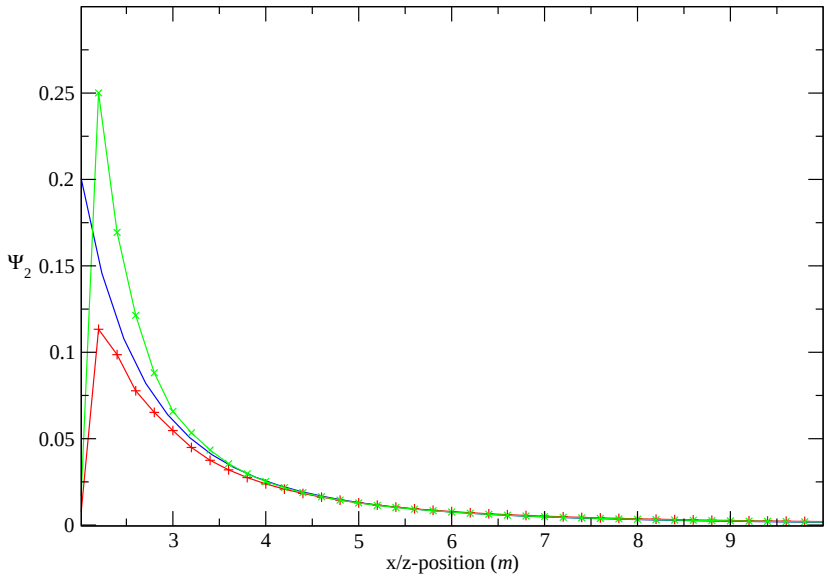

Figure 4.5: The Weyl scalar Ψ_2 along the x-axis (red +) and z-axis (green $\times$), with a fit to the z-axis data (blue), implying $M = 1.63m$.

yields a total mass of $M_x = 1.55m$, while the same fit for z-axis data yields $M_z = 1.63m$.

4.3.3 Bulk Wave Forms

Apart from the bulk mass parameter, we would also like to extract information about the gravitational radiation produced in the on-top interaction. The principal measure of gravitational radiation used here is the "outgoing" Weyl scalar, Ψ_4. This is closely related to the Teukolsky function of BH perturbation theory [127, 128].

In general, the ring-down radiation should be dominated by a spectrum of quasi-normal modes (QNM), as described first by [105]; the relative strength of the modes depends on the precise nature of the interaction that led to the formation of the perturbed black hole. However, in general the first characteristic frequency – the fundamental Schwarzschild $\ell = 2$ mode, whose period is $\approx 16.8M$ (see, e.g., [90]) – should dominate the spectrum (see Section 3.8).

Fig 4.6 shows Ψ_4 as seen by a detector placed at a particular point on the x-axis, near the outer boundary. Strong oscillation is clearly visible to begin with, followed by a few periods of a damped sinusoid, which we interpret as the quasi-normal ringing of the final

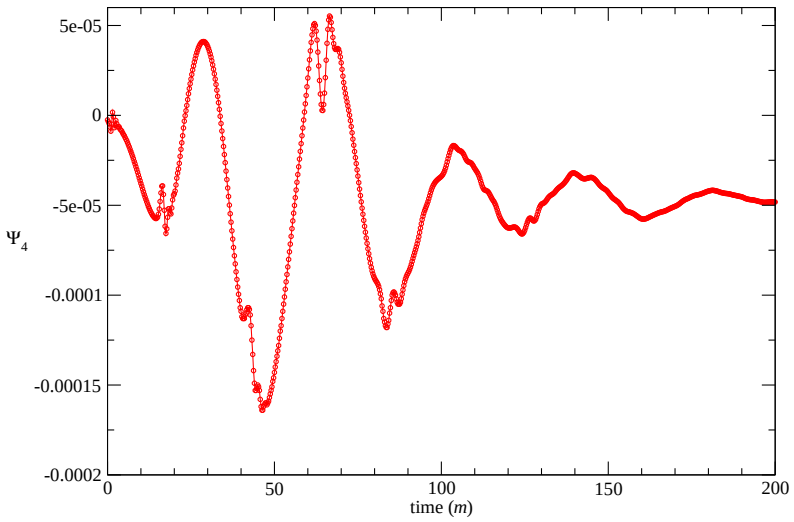

Figure 4.6: The Weyl scalar Ψ_4 at the position $x = 14.0m$ over the first $200m$ of evolution.

black hole. I have taken a discrete Fourier Transform (DFT) of the first $1000m$ of this data (see Fig 4.7). There is a very clear peak corresponding to a period of $T = 37.04m$. Interpreting this as the fundamental $\ell = 2$ mode implies a mass of $M = 2.2m$. Using this base point, I have attempted to match further peaks using Table 3.1. The match can be seen in Table 4.2: there is quite a good match to the first four even-ℓ fundamental modes of Schwarzschild.

Figs 4.8 and 4.9 show an extracted wave form and DFT for the same evolution done in a much larger domain ($30m \times 30m \times 30m$), but with a slightly coarser spatial resolution ($\Delta = 0.25m$). What is immediately apparent is that the spectrum is much cleaner – only two modes seem to be excited significantly. Carrying out a similar matching of the dominant mode yields a background mass of $M = 1.8m$. The other mode is closest to the $\ell = 4$ mode, though the match is not as good as with the smaller domain.

What can be inferred from these two spectra? The two simulations are supposed to be physically equivalent, but the main difference – the domain size – is clearly stimulating modes not affected by the bulk. Outer boundary effects appear both to affect the final BH mass, and excite the underlying hole more.

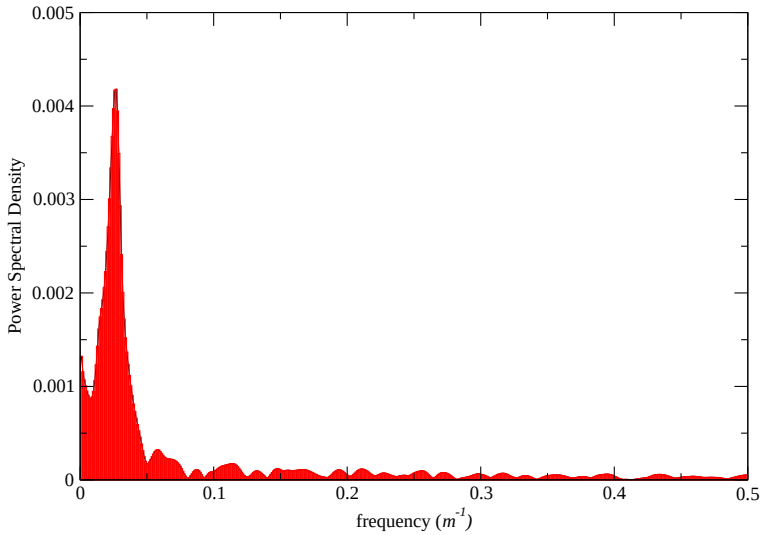

Figure 4.7: Discrete Fourier Transform of Ψ_4 from Fig 4.6 over the first $1000m$ of evolution.

4.3.4 Outer Boundary and Domain Size Effects

The wave form in Fig 4.6 is not particularly clean. This is more obvious if we look at the early stages of the data, shown in Fig 4.10. There is a particularly suspicious pulse passing the detector at $t = m$, and a larger pulse at $t = 16m$. This seems to be pollution from the outer boundary of the computational domain. This is more easily seen in Fig 4.11, where we look at the value of Ψ_4 not at a fixed point on the coordinate axis, but along the whole axis (running vertically) as a function of time (from left to right). The early stages of this plot show (a) a small pulse traveling in from the outer boundary at the initial time step (diagonal from the top left-hand corner of the plot), and (b) a ridge of data appearing after about $15m$ of evolution.

Our interpretation of this pattern is as follows: the ridge is in fact just another pulse, incoming from a transverse direction. It takes around $15m$ for a light-speed mode starting on any of the boundaries to travel to the origin. For points on the x-boundary (that is, all points $x = 15m$), we see a pulse proceeding inward along the x-axis from $t = 0$. For the equivalent pulse beginning on the y-boundary, the first time we see it is after $15m$, when the pulse from every point hits every point on the x-axis simultaneously. This interpretation is illustrated in Fig 4.12, and is justified by a similar run within a newer, irregular domain: Fig 4.13 shows Ψ_4 on the t-y surface when the y-extent is $10m$, but the x-extent is only $7m$. It can be seen now that the x-pulse hits the origin 2-3m before the y-pulse.

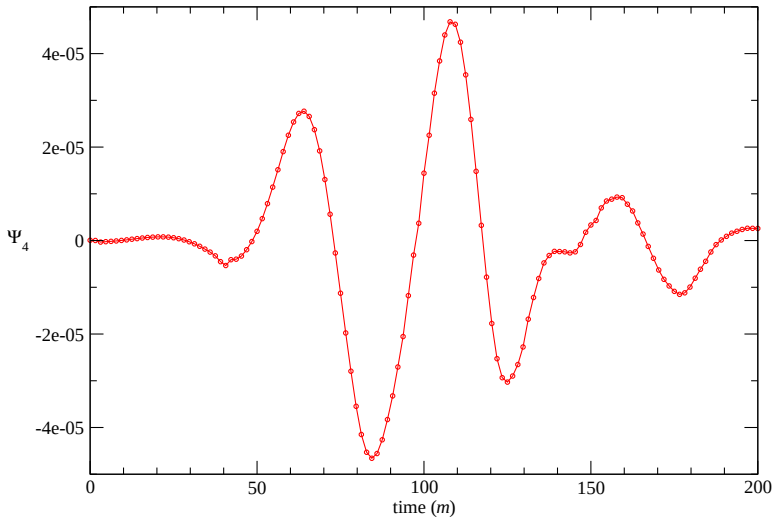

Figure 4.8: Ψ_4 at the position $x = 28.75m$ for a domain of side $30m$.

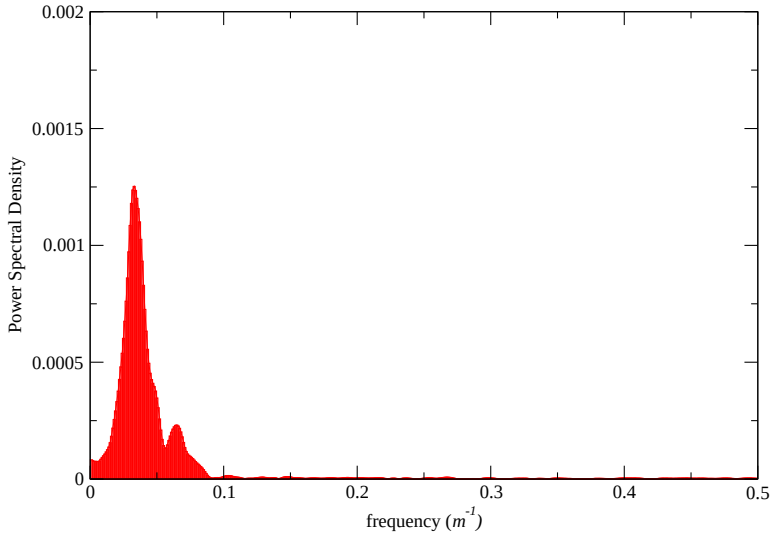

Figure 4.9: Discrete Fourier Transform of Ψ_4 from Fig 4.8 over the first $1000m$ of evolution.

$T_{obs}(m)$	(n,ℓ)	$M_T(m)$	% error	$M_\tau(m)$	% error
37.07	(1,2)	2.21	—	2.03	-8.1
17.26	(1,4)	2.22	0.5	2.07	-9.4
11.51	(1,6)	2.20	-0.5	3.34	51.1
8.78	(1,8)	2.25	1.8	1.91	-13.6
7.58	(1,10)	2.41	9.1	2.68	21.3
30.33	(1,2)	1.80	—	1.69	-6.1
15.40	(1,4)	1.98	10.0	1.81	0.6

Table 4.2: Mode fits to DFTs from Figs 4.7 and 4.9. In each case, the peak position of the (1,2) mode is assumed correct, and errors from other modes and their widths are calculated relative to this.

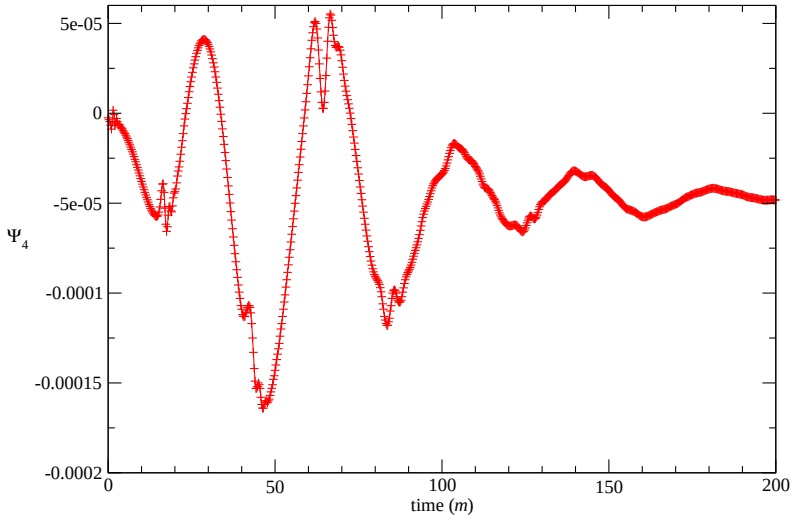

Figure 4.10: Ψ_4 at the position $x = 14.0m$ over the first $25m$ of evolution.

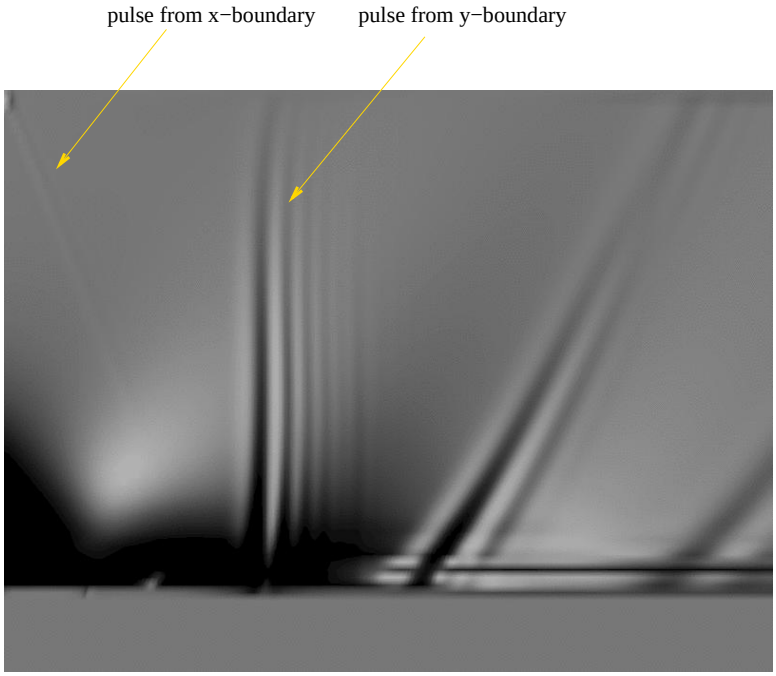

Figure 4.11: $t - x$ surface plot of Ψ_4 over the first $50m$ of evolution. $t = 0, x = 0$ is the bottom-left corner: t increases to the right, while x increases upward.

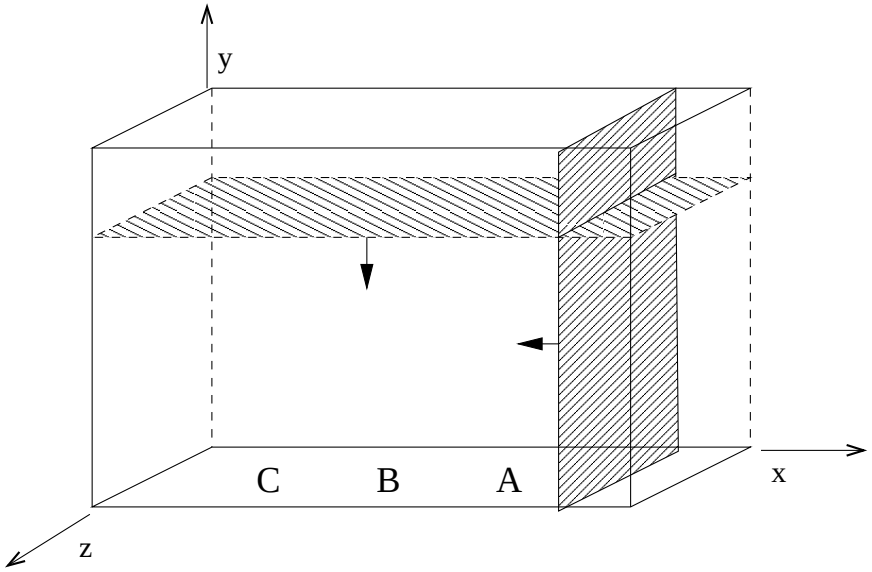

Figure 4.12: Schematic of boundary effects propagating in Fig 4.11. Detectors at positions A, B, and C will see the effects from the right-hand boundary sequentially, but from the top boundary simultaneously.

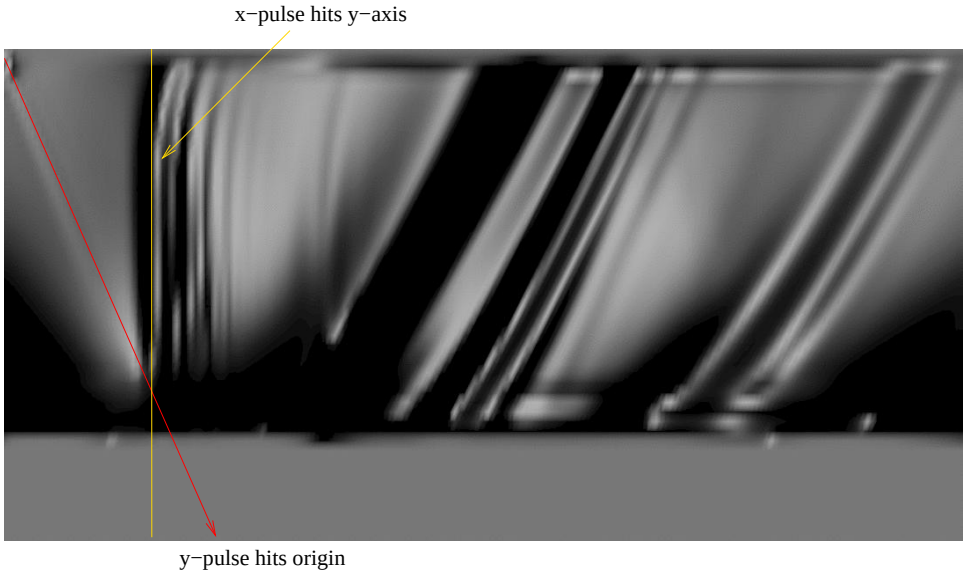

Figure 4.13: $t-y$ surface plot of Ψ_4 over the first $50m$ of evolution for non-cubical domain. The initial pulse from the y-boundary hits the origin at $t \approx 10m$, while the pulse from the x-boundary hits the entire y-axis at $t \approx 7m$.

To investigate this effect further, we compare the waves extracted at the position $x = 11m$ from runs with different domain sizes – $(12m)^3$, $(15m)^3$ and $(18m)^3$. We extract both outgoing radiation Ψ_4 and incoming radiation Ψ_0. Classically, we expect that the Ψ_0 will be very weak for this case – most radiation in the space-time should be outgoing.

Fig 4.14 shows the extracted Ψ_0 for the three domain sizes. Though the wave forms have common features, they are substantially different. Large pulses are seen at times m, $4m$, and $7m$ for the three domain sizes, respectively. These are consistent with the fact that the detector is at a minimal distance m, $4m$, and $7m$, respectively, from the outer boundary faces. A second spike appears at time $13m$ in the smallest domain; a similar spike (of smaller amplitude) occurs for the medium domain at time $16m$, and in the largest domain at time $19m$. Again, the fact that the time between the spikes is equal to the difference in domain sizes is persuasive.

General conclusions we can draw from these observations are as follows:

- The time between the step hitting the three detectors is identical, and consistent with the time taken for a light pulse to travel from the outer boundary.
- The pulse size is a considerable size in comparison with the "background" Ψ_4 itself.
- The pulse is weaker for the larger domains, by the time it hits the detector.

Given the last point, we stick with the largest domain, $(18m)^3$, and look at Ψ_4 extracted

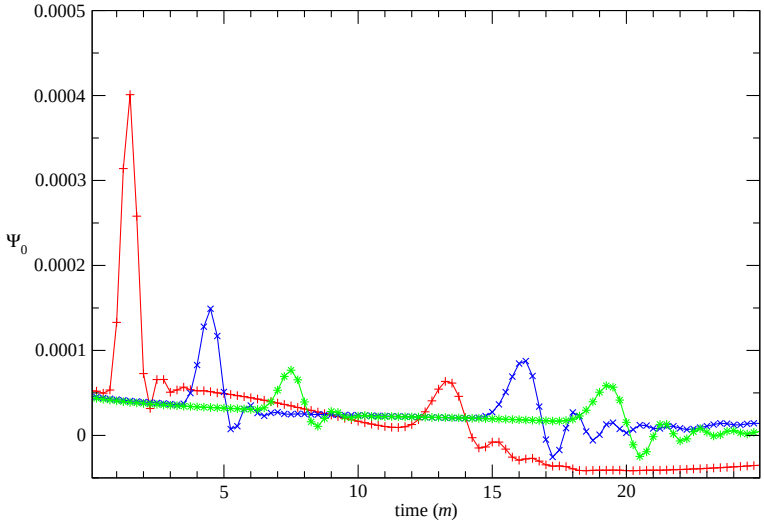

Figure 4.14: Ψ_0 at the position $x = 11.0m$ for domain boundaries at $12m$ (red $+$), $15m$ (blue $\times$), and $18m$ (green $*$).

at three different radii: $11m$, $14m$, and $17m$. This is depicted in Fig 4.15. The most obvious feature from this is the large pulse appearing simultaneously at all detectors along the x-axis at time $t = 18m$. This pulse is the boundary effect traveling in from transverse directions, hitting the x-axis after $18m$ of evolution.

4.3.5 Outer Boundary and Background Mass Effects

Finally, we can look at the dependence of the results here on the mass used for the background "IEF" solution. This solution is subtracted from each of the fields at the outer boundary, and the remnant used as an "error", corresponding to actual outgoing radiation, which is then transmitted out through Sommerfeld conditions. In essence, the background is all we should be left with.

This idea works very well for fields resulting from a run that should actually be stationary to begin with, like the single IEF hole. If the correct IEF mass is used, then the resulting error should be at the level of finite-difference error, and little real radiation will leave the system.

On the other hand, for a non-stationary situation, we really do *not* know what the appropriate mass should be. We assume that for weak perturbations, the energy lost through radiation should be a small correction to the final quiescent BH energy. Say, then, that

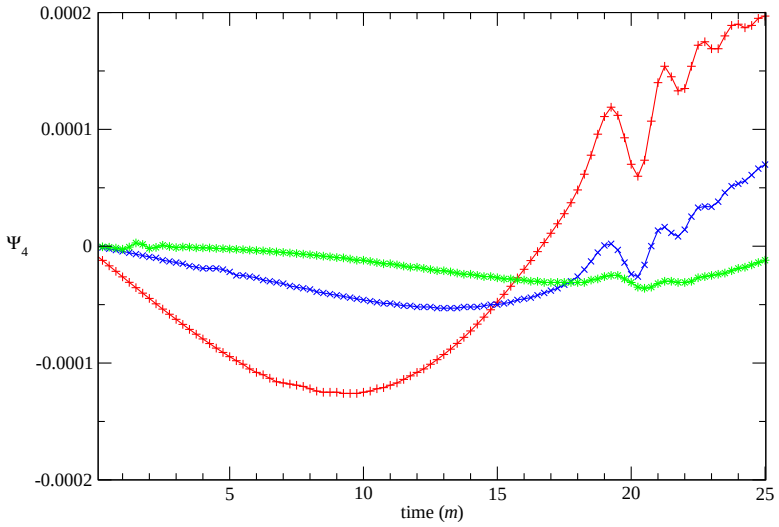

Figure 4.15: Ψ_4 at the positions $x = 11.0m$ (red $+$), $14.0m$ (blue $\times$), and $17.0m$ (green $*$) for domain boundary at $18m$.

we deliberately underestimate the IEF mass for this reason. If we over-underestimate, then we will be left, after subtraction, with an excessively large error. If this is then successfully radiated away by our outer boundary conditions, we will be left with a deficit in the background at the boundary.

One might expect that the initial "pulse" transmitted from the outer boundary would be determined by the error inherent in an unwise choice of background mass. To investigate this, several runs were done with identical parameters (including domain size), but different IEF background mass – $M_b = m$, $2m$, and $3m$. Fig 4.16 shows the early Ψ_0 for these runs; surprisingly, the appearance and amplitude of the boundary pulses seems to depend not at all on the background mass over the large range of masses considered. The most noticeable effect is a growing phase difference in the wave form over time. This leads us to conclude that whatever the cause of the boundary pulses, the exact form of the background solution is not to blame.

4.4 Outer Boundary and Constraint Violation

Before concluding, I would like to note one more disturbing fact related to the outer boundary effects remarked upon so far. Fig 4.17 shows four snapshots of the Hamiltonian con-

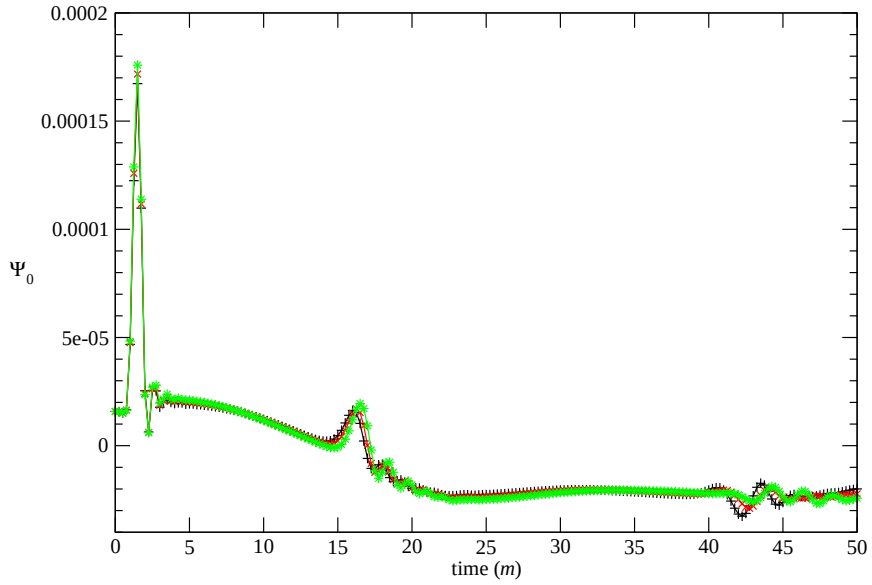

Figure 4.16: Ψ_0 at the position $x = 14.0m$ for background IEF masses of $M_b = m$ (black $+$), $2m$ (red $\times$), and $3m$ (green $*$).

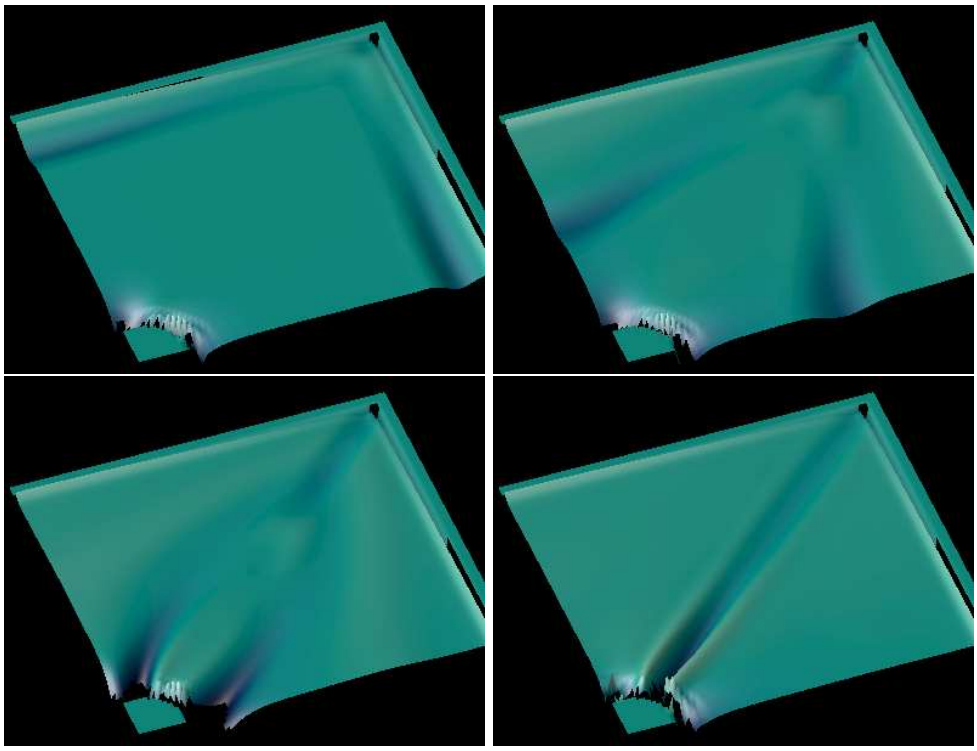

Figure 4.17: The Hamiltonian constraint violation C_H on the x-z-plane at times $t = 10m$, $30m$, $50m$, and $100m$.

straint violation on the x-z-plane at times $t = 10m$, $30m$, $50m$, and $100m$ for the large ($18m$) domain above. The snapshots show the progress of a constraint-violating pulse from the outer boundaries. We are not surprised to observe a pulse of this sort. What *is* surprising is the fact that the violation freezes in at late times: the last snapshot shown is representative of the situation for the remainder of the evolution. In particular, the constraint trench along the diagonal seems generic for all our distorted black-hole runs, as well as for the mergers of the next chapter.

We do not know yet understand the true nature of this frozen-in violation. It is possible that higher-resolution runs of the same physical situation might throw more light on the situation. For instance, if the violation converges away, its unusual shape might be taken to be a consequence of warping in the coordinates near the diagonal. This is quite possible: recall that we have no guarantee that the final coordinates that the system settles down in will be of Kerr-Schild form.

4.5 Discussion

Before moving on to the full head-on merger, what has been learned from this small-separation exercise? The first thing to note is the achievement in evolving highly distorted black-hole data to the point of long-term stability. Our diagnostics indicate that we have indeed reached the space-time of a single black hole in some horizon-penetrating coordinates.

Further than this, we have tried to extract information about the bulk parameters of the space-time – that is, the final mass, M – and the emitted gravitational wave forms. The results of these analyses have been inconsistent and inconclusive, especially in the matter of the total mass of the system. The chief culprit appears to be the nearness of the outer boundary, and the radiative conditions used there.

If anything, the small-separation case is a much less violent one than the full head-on evolutions. The results of this section may be taken to be an upper bound on the quality of the data that we may hope to extract in the next chapter.

Chapter 5

Large Initial Separation: Head-On Collisions

In Chapter 4 I dealt with the evolution of small-initial-separation data, which, since a single well-defined apparent horizon exists on the initial time slice, we can treat as the ring-down of a single perturbed black hole. In this chapter, I move on to the collision of holes where no common apparent horizon can be found initially[1].

It should be pointed out that although full 3D merger simulations have been carried out – with varying levels of success – since 2000, almost all of the work done has been with Bowen-York-type data and puncture-type evolutions. The present work is the only significant work done with both Kerr-Schild-type data and moving singularity excision.

Of course, the equal-mass head-on problem is a significantly easier one to tackle in general than a grazing, unequal-mass, spinning black hole interaction. However, with the single exception of making use of octant symmetry for speed and memory requirements, we have not exploited this simplicity. Our evolutions are fully 3D with the full set of Einstein's equations.

In Section 5.1, I briefly review the history of the BBH collision problem, and how it has been approached numerically since the 1960s.

Since we are dealing with two moving holes, we need a strategy for dealing with the slicing and boundary conditions; these are addressed in section 5.2.

As the holes will initially have a large separation, we will need to excise separately within the apparent horizon of each hole; that is, there will be two separate excision masks. Moreover, these masks will naturally have to move as the singularities converge in the center of the domain. In section 5.3, I discuss the problems inherent in tracking black holes moving in the computational domain, and the procedures introduced in the Maya code.

At some stage, the individual masks will give way to a single larger excision mask in the center of the domain. In section 5.4, I present the problems inherent in switching to a merger-hole regime, and how these were addressed.

In section 5.5, I present full merger runs, with a well-defined regime switch from two well-separated black holes to a single perturbed merged hole.

[1] As discussed in Section 3.5, absence of a common AH does not prove that both holes are not within a single event horizon.

Finally, section 5.6 discusses these results, and outlines extensions to this work and future directions.

5.1 History of BBH Mergers

The history of the equal-mass head-on collision problem is almost the history of vacuum Numerical Relativity itself. The problem possesses natural axial symmetry, and was practically the first to be addressed on digital computers. The first data sets considered were Misner data [93, 94] – conformally flat, with a two-sheeted topology. N holes are represented by N Einstein-Rosen bridges between our external universe and another internal one.

Numerical evolutions in 2D were attempted as early as 1964 [79]. These, and subsequent evolutions over the next decade-and-a-half, were done in axisymmetry, and almost always with Misner data. A major push by DeWitt and collaborators in the 1970s led to actual wave forms for Misner-type data in fairly close-separation simulations. These results are summarized by Smarr [119].

While elegant, the Misner two-sheeted topology necessitates an infinite method-of-images summation to satisfy consistency conditions at each throat. Brill and Lindquist [33] data sets are similar in many ways to Misner data – they're conformally flat initial data, with the singularities hidden inside the conformal factors. The difference is in the topology – N holes for Brill-Lindquist means $N + 1$ sheets. The "upper" sheet, where all N throats terminate, represents our universe, while the remaining N sheets are completely disconnected, each representing the asymptotically flat space near that hole's singularity. Though slightly less elegant than Misner data, it needs to satisfy no consistency conditions.

Later work in the early 1990s by Anninos et al cemented the 1970s results, with more accurate wave forms. By now, computer speed and memory had improved to the point that fully 3D evolutions could be attempted. Bowen-York data [28], an extension of Brill-Lindquist that allowed for spinning and boosted black holes, became popular, and further head-on and grazing collisions were performed.

The first fully 3D evolution of a binary system was by Brügmann in 1999 [35]. This used the new BSSN formulation of Einstein's equations, and adaptive mesh refinement (AMR) to resolve the high-curvature region. The new "puncture" method was used to remove the physical singularities from the evolution domain.

In 2000, Brandt et al [31] set out to evolve a different kind of initial data – the so-called "HuMaSh" data [92], a superposition of the spatial parts of two Kerr-Schild solutions. Their simulations were of grazing collisions, with separately spinning holes initially boosted obliquely toward each other. The physical singularities were treated using singularity excision. These grazing collisions were apparently successful, but the simulations died after

about $15M$ of evolution. The following year, Alcubierre et al [6] carried out similar simulations for Bowen-York puncture data, with runs lasting for roughly $30M$.

Also in 2000, the "Lazarus" project was launched [16]. Given the limited evolution times available for mergers, this project aimed to take the data generated from fully 3D merger runs before those runs died, and use it as a starting point for a perturbative evolution. This scheme requires that we know that the system is in the perturbative regime before the end of the 3D evolution, and that the data at that stage is dependable.

Since the advent of puncture initial data and the wide acceptance of BSSN and its variants as good evolution systems, much recent work in full 3D Numerical Relativity has focused on generating quasi-stationary binary BH orbits. The ability to evolve a binary system for several periods of a (nearly) circular orbit before plunge and merger would essentially complete the binary black hole program: far-separated holes could be analyzed via the post-Newtonian (PN) approximation, taking us close to the innermost stable circular orbit (ISCO). From then on, full numerical evolution allows us to probe the strong-field gravitational interactions of merger. Post-merger, if necessary, perturbative methods can take over, evolving the nearly stationary space-time until it achieves a final steady state.

Up to the end of 2003, the best orbit simulations were struggling to manage a fraction of a circle. Brügmann et al [36] have reported significant progress in this, and may have achieved an orbit.

Thus, most of the current work on binary systems has used conformal (Misner, Brill-Lindquist, Bowen-York) data, and is approaching dynamical situations like binary orbits using quasi-stationary tricks. Much has been achieved, and continues to be achieved, through this "motion-without-motion" approach. Nevertheless, it is my opinion that such restrictions have a limited shelf life: sooner or later, even carefully crafted initial data sets and slicing conditions will need the ability to move singularities around the numerical grid.

For this reason, we have returned to the head-on merger using the same approach as Brandt et al. We are using HuMaSh data – ideally suited to singularity excision – and testing what the intervening years' progress have bought us. We aim to improve and extend their results with the extra insights and tools that have become available since then.

5.2 Slicing and Outer Boundaries

Once again, given initial data, this time representing two well-defined black holes, we need to settle two main issues before we can begin evolution: boundaries and slicing.

The situation *vis-à-vis* outer boundaries is essentially the same as for the small-separation case: to use radiative boundary conditions on the outer boundary, we must first choose a background metric. The only difference here is how much less suitable single IEF is to represent the background. For the most part, in what follows, we use the unsolved binary

HuMaSh as the background itself. Although certainly not ideal, it should have roughly the correct fall-off rates for the various fields. I shall justify this numerically later.

The slicing situation, by contrast, is much more serious than in the small-separation case. Then we were already comfortably inside a common horizon, and could use the 1 + log and gamma-driver conditions to allow the lapse and shift to settle down to a steady-state form. In general, however, there is no known live gauge condition that would allow us to evolve the lapse and shift from their initial data, when not close to a perturbed single black hole. In fact, I am not aware of any such live condition even for a single boosted hole moving across the evolution domain (where we know the analytic solution).

This means that we are again forced to make use of some analytic background to determine an appropriate slicing condition. In principle, we could use any lapse and shift we wished – after all, they are gauge variables. Unfortunately, as early experiences with geodesic slicing show, the most obvious slicing schemes may cause the evolution to crash from purely coordinate effects. Thus the most obvious analytic slicing condition is probably one generated in a similar manner to the binary data itself.

Thus we find that, for the time being at least, we need to use analytic binary data for peripheral purposes – outer boundary background and slicing – and so we also need the appropriate parameters for such data: masses, singularity positions, and boost velocities.

5.3 Excision and Singularity Tracking

5.3.1 The Need for Gaussian Tracking

For a dynamic BH simulation, with one or more BHs moving across the computational domain, an excision-based code (such as Maya) will need to move the excision mask accordingly, to ensure that the BH singularities are safely hidden from the computer. In Chapter 2, I introduced the dynamic excision algorithm used in the Maya code, with its extrapolation and repopulation procedures. I also mentioned the GaussTracker module that has allowed Maya to move the excision mask naturally over the grid for circling and spiraling coordinate motion.

For the more physically motivated problems discussed here, the GaussTracker is an attractive mask update scheme, for two basic reasons. First, the use of an automated mask moving method is aesthetically very pleasing – what we want in the future is to be able to bring NR to the point of a real initial-value problem, with no user intervention after $t = 0$. The GaussTracker module is an encouraging sign that such an aim is achievable.

Secondly, and more practically, we need this scheme because a simple use of the singularity position would not lead to a very well-centered excision mask for boosted binary data such as the sort we consider here. For pure coordinate motion, K will have a center

of curvature more-or-less exactly co-incident with the actual singularity position. However, for situations involving boosted data, this will no longer be true. Take, for instance, K for a single IEF hole boosted with a velocity v in the positive-z direction. The K here will then deviate from the non-boosted form. For small boosts, we can approximate the full K to linear order (see Appendix, Section A.3):

$$K \approx \frac{2M(r+3M)}{\sqrt{r^3(r+2M)^3}} - v\frac{Mz(r+8M)}{\sqrt{r^3(r+2M)^5}} \tag{5.1}$$

Figure 5.1 shows the full profile of K along the z-axis for a significant boost speed, $v = 0.7$; overlaid are the GaussTracker-determined centers at several different excision radii, their deviations from the actual singularity position (the origin) indicated by line segments of differing lengths. We can see here that the center chosen by the GaussTracker procedure will depend strongly on the excision radius chosen. The smaller the radius, the better the mask center will approximate the true singularity position. What we must weigh against this is the greater gradients we encounter as we approach the singularity more closely. Moreover, these gradients make the extrapolation procedure more risky. Thus we are restricted in how well the GaussTracker can actually perform in estimating the singularity center.

A similar effect is seen when we are dealing with data describing an unboosted binary BH system – the influence of one hole will affect all fields near the other singularity; the difference now is that we have no analytic knowledge of where the singularity is after the initial step. This fact makes the output of the Gaussian fit method a very bad estimator of the actual singularity position, except in the limit of an arbitrarily small excision radius, but an excellent pragmatic fit for the excision mask itself, whatever the situation. This is because for excision, the actual location of the physical singularity is less important than where steep gradients lie. It's better to minimize the gradients than to aim always to center on the singularity.

It also seems to be a good idea, in highly dynamical situations, to avoid using the excision boundary points in the Gaussian fit. The rationale is simply as follows: in very dynamical situations, the singularity, and hence the associated gradients, will move quite rapidly; therefore, it is likely that excision boundary points will only have a few time steps to settle down after data population before being used to recenter the mask again. Although it helps to set a limit on how often the Gaussian algorithm is invoked (or rather, how often its results are actually used to recenter the mask), the excision boundary data is still inherently less reliable than that even one data point away. This observation will be justified experimentally later.

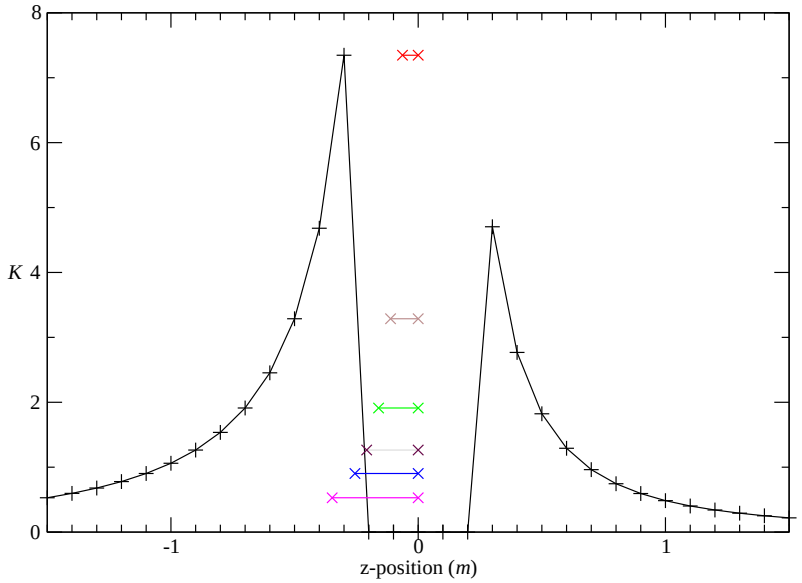

Figure 5.1: K for highly boosted ($v = 0.7$) IEF data, excised at an ideal radius of $0.4m$. The horizontal line segments correspond, from top to bottom, to excision radii of $0.4m$, $0.6m$, $0.8m$, $1.0m$, $1.2m$, and $1.6m$: in each case, the end points are the true singularity position (that is, the origin), and the center determined by the GaussTracker algorithm.

5.3.2 The Need for Singularity Trajectories

I argued above that we cannot rely on the singularity positions to center the excision masks well. This argument works equally well in reverse: if we have obtained good mask centers (from, for instance, GaussTracker), we cannot use them to locate the singularities, due to the asymmetry of the field K on which the mask determination relied. As we decrease the excision radius, the mask center becomes a better approximation to the singularity position, but the increasingly steep gradients present as a result would quickly kill the simulation.

One might also consider locating the singularity using a more invariant diagnostic. The apparent horizon, for instance, is a strong indicator of the presence of a black hole. Brandt et al used the horizon center to determine a rough trajectory for the pre-merger black holes in the grazing collision. We must consider, however, the warping that can happen to an apparent horizon for different situations. The horizon for a single Schwarzschild black hole will be a spherical surface at a (Schwarzschild) coordinate radius of $r_{AH} = 2M$. For a spinning Kerr hole, the horizon will still be a sphere, but with a smaller radius, $r_{AH} = M + \sqrt{M^2 - a^2}$. For a boosted Schwarzschild hole, the horizon will be length-contracted along the boost direction, when viewed in the lab frame. Though there has been a change of shape, the geometrical center of the horizon is still at the physical singularity. For small boosts, to lowest order in v,

$$r_{AH} = 2M + Mv^2 \left[-\frac{104}{2107} - \frac{2159}{2107}\cos^2\theta - \frac{13}{903}\cos^4\theta + \frac{1}{43}\cos^6\theta \right]. \tag{5.2}$$

It is even possible that the geometrical center of the apparent horizon will be a good estimate in more complex situations, for instance, in the case of binary data. Bonning et al [27] show that, for HuMaSh data with widely separated singularities, the horizon for hole A (centered at the origin) will be located at points satisfying

$$\begin{aligned} r_{AH} &= 2m_A + \left(\frac{m_A m_B}{3d} P_0(\cos\theta) - \frac{m_A m_B}{2d} P_2(\cos\theta) \right) \\ &= m_A \left(2 + \frac{m_B}{12d}(7 - 9\cos^2\theta) \right). \end{aligned} \tag{5.3}$$

This again describes a length-contracted coordinate shape – the two extremes in the z-direction are at $\theta = 0$ and $\pi/2$. However, in the strong-field interaction when the holes are closer together, we cannot expect such symmetry to persist.

Another issue is the numerical accuracy of the horizon finder used. Depending on the algorithm employed, the solver might gain in error when the hole is not on a symmetry axis. In a case such as this, the geometrical center may be hard to calculate accurately, and even hard to define unambiguously.

It should be mentioned here that though Brandt et al [31] did use an apparent horizon

finder to determine both singularity and mask centers in their grazing collision runs, this was not a full feedback procedure. They found apparent horizons until the code crashed due the singularity being unmasked, calculated a trajectory from the horizon centers so far, and used this trajectory to evolve the centers a bit further. Additionally, since their evolutions were of a fairly short duration, with a limited separation, the error involved in such an approximation would not have been fatal.

This is the state of affairs at the time of writing. It is unfortunately the case that we are still limited to constructing a trajectory for the BH singularities. The trajectories used are "pseudo-Newtonian" – the initial positions and velocities are consistent with those of two particles falling in from rest at spatial infinity, but the motion thereafter is simply that of uniformly accelerated particles. This was a choice of convenience at the time, but it seems to work quite well in practice. It is true that the acceleration may have to be adjusted to ensure a long-running evolution before merger; an additional "jerk" parameter – the constant time-derivative of the acceleration – was introduced. Then the singularity position would be, at a certain coordinate time t:

$$\vec{r}_c(t) = \vec{r}_0 + \vec{v}_0 t + \frac{1}{2}\vec{a}_0 t^2 + \frac{1}{6}\vec{j}_0 t^3 \tag{5.4}$$

A non-zero "jerk" was not always needed; many of our mergers worked with constant acceleration.

5.4 How to Merge Two Black Holes

So far, I have addressed the type of initial data used for a head-on merger run, and the manner in which the singularities and excision masks are treated during binary evolution. The post-merger, ring-down phase of evolution will be similar to what was described for the close limit in Chapter 4. All that remains is to describe the merger procedure itself.

5.4.1 The Merging of Excision Masks

Of course, the behavior of the physical evolution fields is entirely deterministic, and nothing needs to be done to persuade two incoming black holes to merge into a larger one. However, singularity excision is a technical construct, and we *do* have to intervene to achieve the merger of two initially distinct excision regions into one common region surrounding the forming singularity of the end-state merger hole. Pseudo-code is shown for the mask merger procedure in Fig 5.2. The code is run once every time step.

We must be careful when deciding to excise a new region from the computational grid. Our justification for excising originally was that the excised points were comfortably contained inside a region causally disconnected from the outside world. We establish this

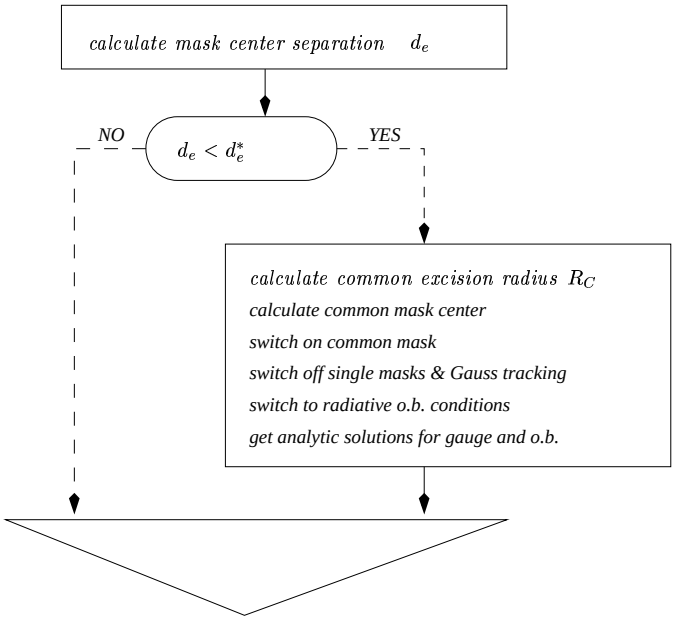

Figure 5.2: Flowchart for merging the excision masks.

disconnectedness by verifying the presence of an horizon surrounding the excised points. Thus when switching to a new common excision region, we first demand that a common horizon is present, surrounding both holes. Additionally, to preserve our causal disconnection, our new excision region should certainly also be contained fully by the common horizon.

When evolving binary data with any large initial separation, the first common horizon to be found will not look very spherical in our numerical coordinates: distortion will tend to give it a peanut-like shape, possibly even locally concave at some points. This horizon will generally only just contain each of the single-hole horizons. A coordinate sphere centered at the origin and containing each of the individual excision regions will certainly not be contained in this early common horizon. Unfortunately, the present version of the Maya code contains the facility to form only fairly simple excision shapes – cubical, LEGO-spherical, and LEGO-cuboctahedral. None of these shapes will be able to accommodate the two individual excision regions while simultaneously staying inside the new common horizon.

This is a temporary obstacle: over time, the individual holes will get closer still, and the common horizon will circularize considerably. If we wait long enough, we can comfortably excise a coordinate sphere without violating causality. Fig 5.3 outlines the necessary logical stages of this process. However, we cannot wait indefinitely before merging; in fact we cannot even wait until the two excision regions are touching. As discussed in Sec 2.7, we *must*

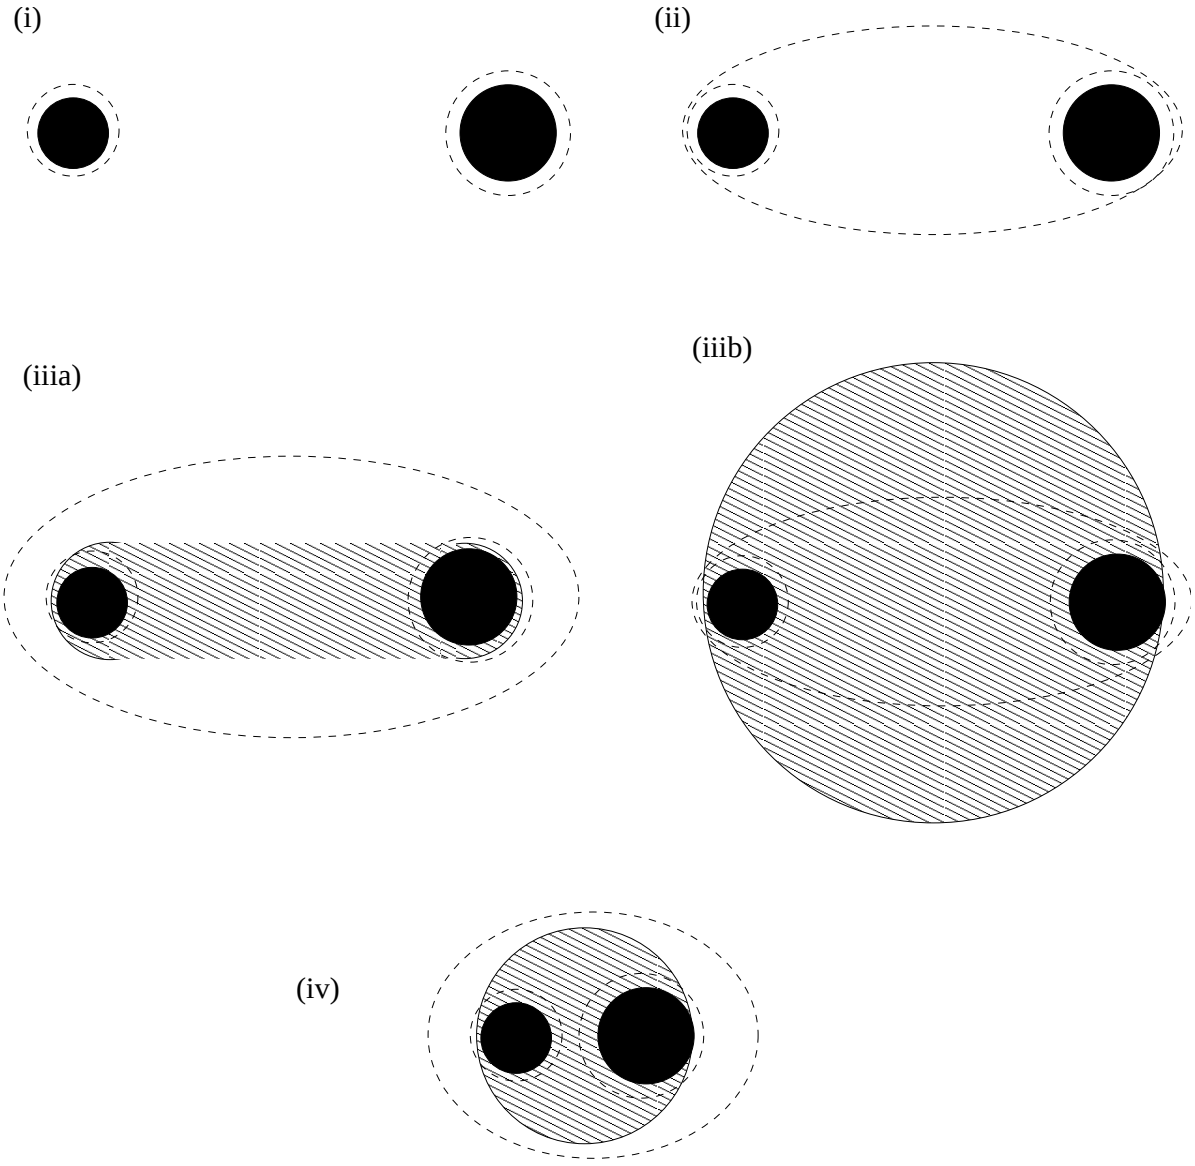

Figure 5.3: Necessary conditions for a merging of excision masks during BBH head-on collision. Dashed lines indicate marginally trapped surfaces, black regions are already-excised points, while shaded regions are proposed common excision masks.(i) No common apparent horizon present; (ii) First common apparent horizon detected; (iiia) The ideal common excision region will be as efficient as possible, just containing previously excised points; (iiib) The real common excision mask must be a coordinate sphere; (iv) To get a spherical mask while respecting causality, we must wait until the singularities have converged more, and the common horizon has circularized.

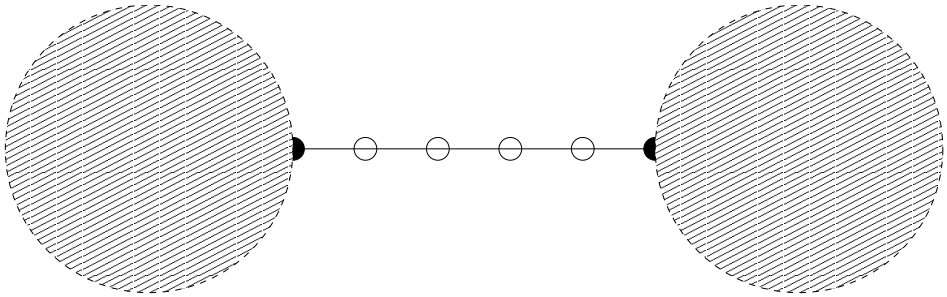

Figure 5.4: The pre-merger excision masks shortly before merger. There must be at least k_{xpol} in-domain grid points – denoted by open circles – available for extrapolation at all times.

have k_{xpol} points available at all times for extrapolation procedure (Fig 5.4). Depending on the spatial resolution used, this may be impossible to satisfy at the same time as the causality condition.

Ignore for the moment the imperfect LEGO nature of the excision boundaries. Take a point on the surface of one of them on the line joining the two centers. This point is an excision boundary point, and field values here must be filled in by extrapolating using data from interior (evolution) points along the line.[2] If there are n interior points along this minimal distance line, then at the closest point, the edges are separated by a distance of $(n+1)\Delta$.

Now consider the actual shape of the excision region. This will be determined by an excision radius R. For a spherical excision shape, the distance from excision center to excision center is then

$$d = R_A + R_B + (n+1)\Delta \tag{5.5}$$

This is a lower bound on the separation of the excision regions. For spherical excision post-merger, this d will be related to the radius of the new common excision region, which has to fit wholly inside the common horizon:

$$R_C = \frac{1}{2}\left(R_A + R_B + d\right) = R_A + R_B + \frac{(n+1)\Delta}{2} \tag{5.6}$$

At low spatial resolution, the Δ factor can be a serious impediment to fitting the new mask inside the common horizon. At the resolution used for most of the simulations here, $\Delta = 0.2m$, two single-hole masks of radius $0.8m$ imply a total excision radius of $R_C = 2.1m$. That is, single-hole masks with 40% of the single-hole horizon radius mean the new excision radius is approximately 51% of the common horizon radius (assuming a total mass of $2m$).

[2] Of course, when one is using points near hole B to extrapolate onto hole A's excision boundary, the results may not be meaningful, but we are looking only at the technical requirements.

This is essentially the reason why we cannot attempt mergers with too coarse a resolution: the extra radius due to the grid spacing will push us through the threshold. In practice, even a resolution of $\Delta = 0.25m$ has crashed runs at the time of mask merger.

5.4.2 Switching to Ring-Down Slicing and Boundary Conditions

There is one more area that has to be addressed before the merger discussion is exhausted: how to switch to a long-term stable live slicing condition post-merger. Fig 5.5 is a flowchart for the procedure followed to merge the singularities for gauge (and outer boundary) purposes.

Single-hole evolutions do much better when the slicing condition is live, rather than algebraic; we would like to take advantage of this to achieve as much stability as possible in our system post-merger. Results from Chapter 4 indicate that this is possible.

However, as noted above, we are forced to evolve until merger with analytic α and β^i, given by some appropriate binary data, such as (3.19-3.20), or (3.22-3.23). Our general strategy is that by the time we switch the masks, we are "close" to a single-hole solution, and then can switch to live slicing.

What freedom, then, do we have in our slicing switchover? One general choice is the exact timing of the switchover. The plausible window for this runs from the moment we merge the masks to the point at which the pseudo-Newtonian trajectories of the two holes meet at the center of mass. Our initial simulations used the very start of this window for switchover, but a possible reason to defer the slicing switch until later is the shape of the lapse and shift at switchover – closer singularities mean a slicing that is closer to the single Kerr-Schild limit. Against this we must balance the increased anisotropy gained by having larger boosts.

Given that the lapse and shift will not look perfectly like a single Kerr-Schild lapse and shift, we might also decide to restart the lapse and shift as clean single-hole quantities just as they go live. This may appear to be cheating, but lapse and shift are gauge quantities. We are free to do with them as we please, as long as the simulation stays stable.

Another possibility considered was the gradual blending of slicings from binary slicing (up to the point of mask merger) to single-Kerr-Schild (at singularity trajectory merger). Then we could evolve freely from that point on, without worrying about the recalculation of the lapse and shift.

The simulations presented in the next section were, for the most part, performed with lapse and shift being recalculated once at the center of mass, and allowed to evolve live from then on.

It should be emphasized that the justification behind these different schemes is *not* to ensure stability – the evolutions are long-term stable whichever choice we make – but to try to cut down on gauge noise in evolved quantities, and thus improve the extracted wave

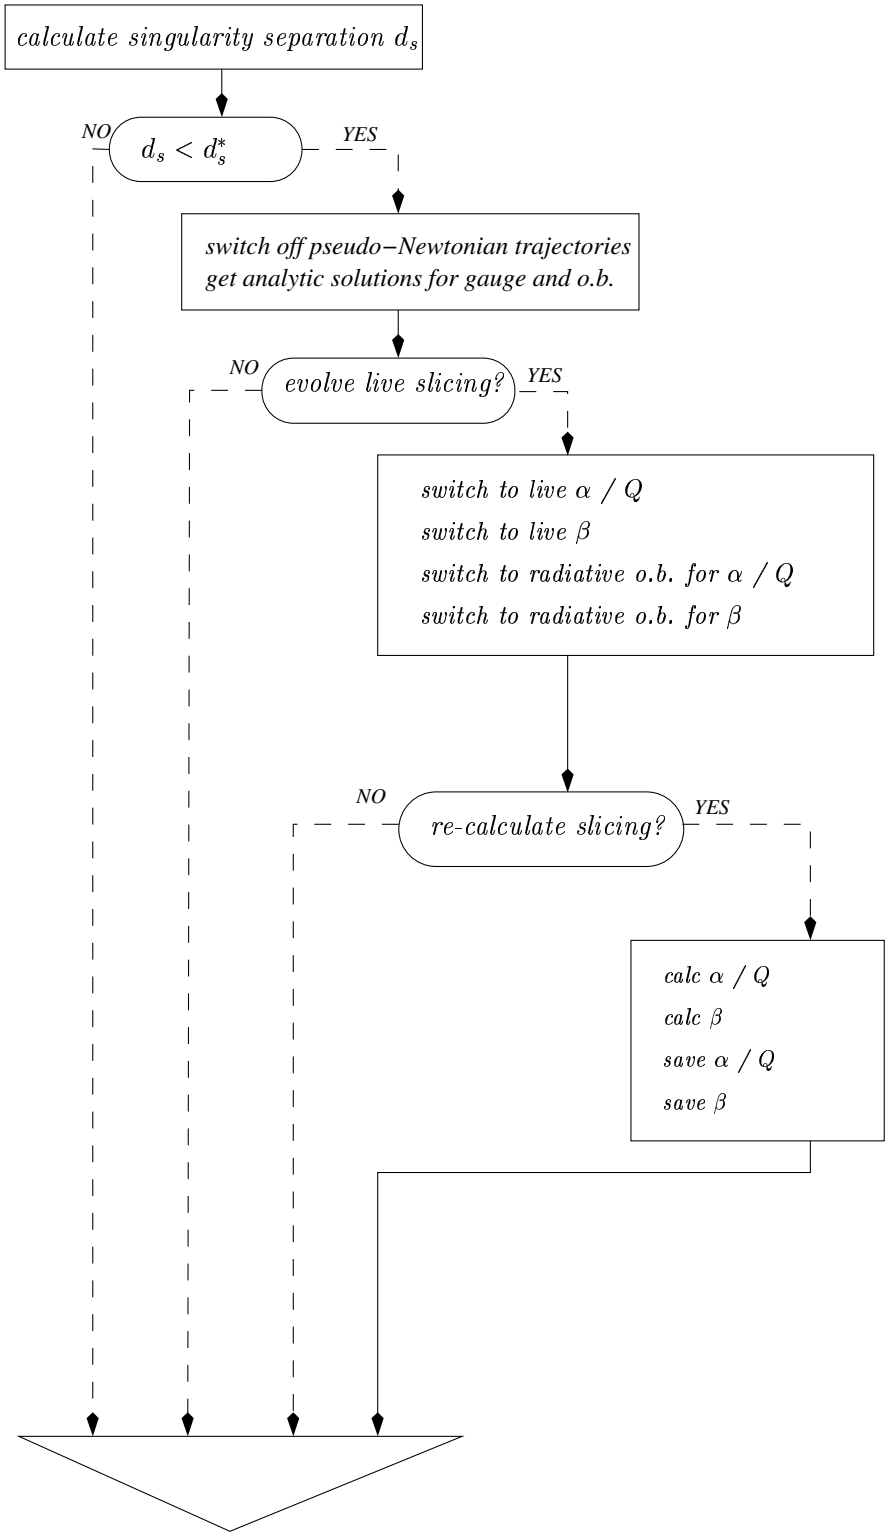

Figure 5.5: Flowchart for merging the singularities.

forms.

5.5 Full Merger Simulations

I now present sample results of this collision-and-merger procedure. Initial data is the same HuMaSh-type boosted, non-spinning KS binary data. We have performed runs for a range of coordinate separations – $d = 12m$, $16m$, and $32m$. Details of the parameters of each run can be found in Appendix C.

What all runs have in common is their spatial resolution, $\Delta = 0.2m$. They also share the same Courant factor, 0.25, implying a time spacing of $0.05m$. As discussed above, coarser spatial resolution will often prevent successful excision mask merger. Finer resolutions, however, mean a very limited physical domain size. We have taken $0.2m$ as our compromise between these two issues. This means, however, that convergence tests have not been practicable up to this point.

For memory efficiency, most of these mergers were carried out in octant symmetry. This might seem strange, given that we're discussing a binary system, but the reflection symmetry of these problems means we lose no information in doing so. In octant symmetry, allowing for inter-processor "ghost zones", a domain of size $(18m)^3$ means $(96)^3$ grid points in total. This is not a huge load – a single modern PC processor will have the speed and memory to handle more than 50^3 points on its own. The real bottleneck here is the elliptic solver used for the initial data. The method of solution is very memory-intensive (as discussed in Section 3.3), and will begin using disk swap space for runs of larger than $(120)^3$, no matter how many processors we throw at the problem.

Our diagnostic of the quality of merger, the apparent horizon finder, was run once every m in time (that is, every 20 time steps).

5.5.1 Stability

To begin with, I will demonstrate that we can, indeed, follow two well-separated black holes to the point of merger, and additionally, that the merger procedure and more widely separated data do not destroy the long-term stability observed for the on-top runs. We carried out a long simulation for the simplest merger procedure, using the parameters of the $16m$-separation run from Section C.2.

In this run, the mask merger was triggered at $t = 17.05m$, and the gauge merger happened at $t = 18.75m$. At the latter time, the densitized lapse Q and shift β^i were recalculated from the analytic solution for a single IEF hole of mass $M = 2.0m$, centered at the origin. From this time on, all variables were evolved (the lapse and shift according to the 1+log and gamma-driver conditions, respectively), with radiative boundary conditions on an IEF background.

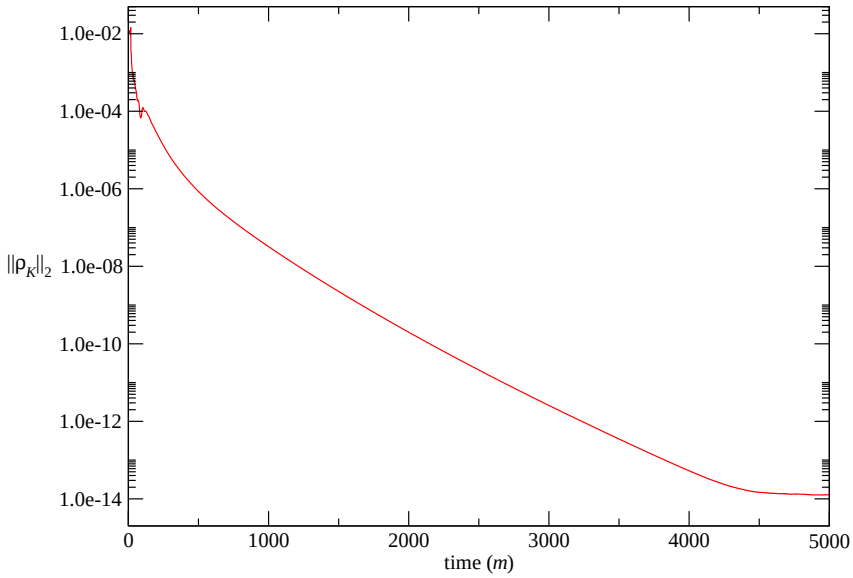

Figure 5.6: Stability of full head-on mergers for initial separation of $16m$.

Fig 5.6 shows a semi-log plot of the $L2$-norm of the ρ_K, the right-hand-side of (2.21). As can be seen, K begins to settle down exponentially after about $500m$. We evolved until $t = 5000m$, at which point the variation had leveled off at just over 10^{-14}.

As with the small-separation runs, we may look at the speciality index S to assess the extent to which the run has settled down to a stationary type-D situation. Fig 5.7 shows S for two times: just post-merger, and near the end of the simulation. We can see that at late times, the solution is certainly of type D. There is a small deviation from unity at the excision boundary, but this can be attributed to finite-differencing error.

5.5.2 BH Mass Determination

As with the small-separation case, we would like to determine the bulk parameters of the end-state space-time. Given the type-D nature of the late-time numerical solution, and that the symmetry of the head-on collision precludes net angular momentum, the late-time solution must in fact be a Schwarzschild black hole.

Fig 5.8 shows the ADM energy and apparent horizon area as calculated over the first $1000m$ of the evolution. We notice two things simultaneously: (1) the final ADM mass is very low indeed – $M_{ADM} \approx 1.2m$ at late times; (2) the apparent horizon is consistently not

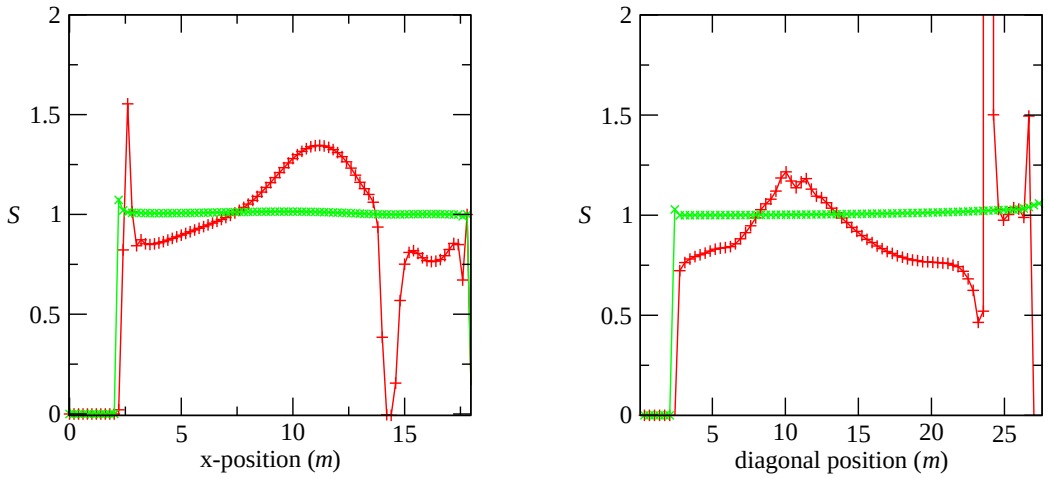

Figure 5.7: Speciality index S just after merger (red +), and at late times (green $\times$). The left plot is along the x-axis, while the right-hand one is along the diagonal.

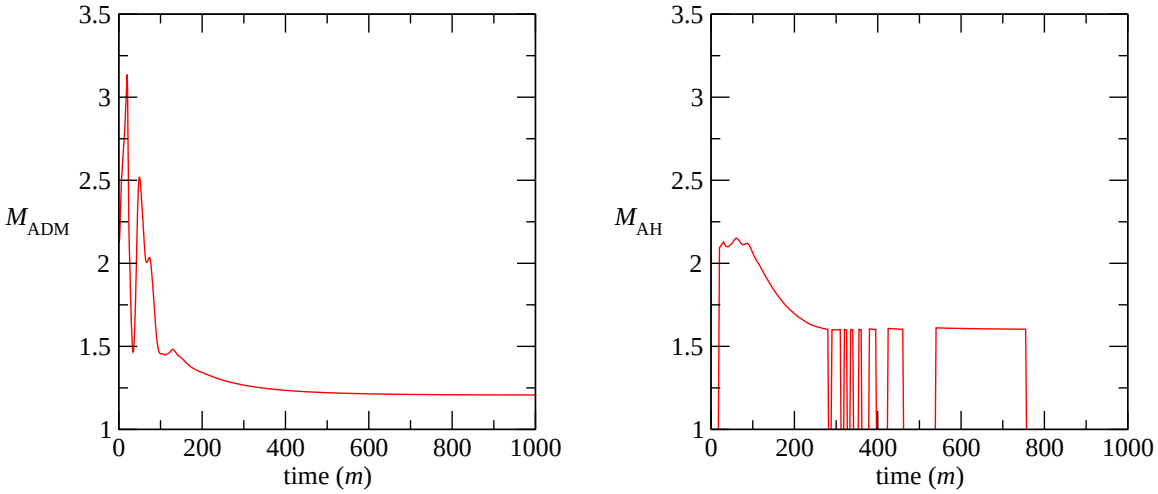

Figure 5.8: The ADM energy (left) and apparent horizon effective mass (right) as a function of time. The failure to locate an horizon is shown by a fall to zero of M_{AH}.

found at late times. This can happen if the horizon finding algorithm looks at points that we have excised from the computational domain. However, our final excision region has a radius of $R_C = 2.5m$. If the horizon finder encounters points within this radius, it implies that the actual horizon must be within a grid spacing or so of the excision radius, lending support to the lowness of the ADM estimate. The last horizon estimates for the mass are around $M_{AH} = 1.61m$.

Assuming the appropriate coordinates are still in place, we should be able to extract the BH mass from the monopole term, Ψ_2. Fig 5.9 shows the late-time ($4000m$) Ψ_2 along the x- and z-axes. A fit to this data yields a very low estimator of the BH mass – between $0.58m$ and $0.87m$. Possible non-physical reasons for this are the lower reliability of derived quantities near the singularity, and a large drift of the radial coordinate from the ideal IEF r.

5.5.3 Bulk Wave Forms

We expect to see two characteristic types of wave forms from these runs: first waves representing the end of the strongly nonlinear merger process; then the characteristic spectrum of a perturbed black hole, which marks the ring-down to a quiescent state.

Fig 5.10 shows Ψ_4 as extracted from the $16m$-separation run, at a fixed point on the x-axis near the outer boundary.

As with the on-top case, we have performed a DFT of the Ψ_4 wave form. The result for an observational window of $t = 100m$ to $1000m$ is the spectrum shown in Fig 5.11. We waited until after $100m$ of radiation to allow the pre-merger and merger radiation to leave

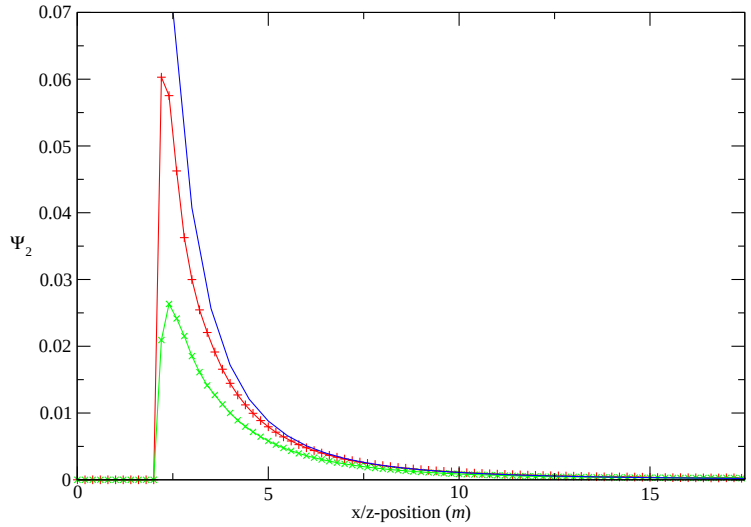

Figure 5.9: The Weyl scalar Ψ_2 along the x-axis (red +) and z-axis (green $\times$), with a fit to the x-axis data (blue), which implies $M = 0.87m$.

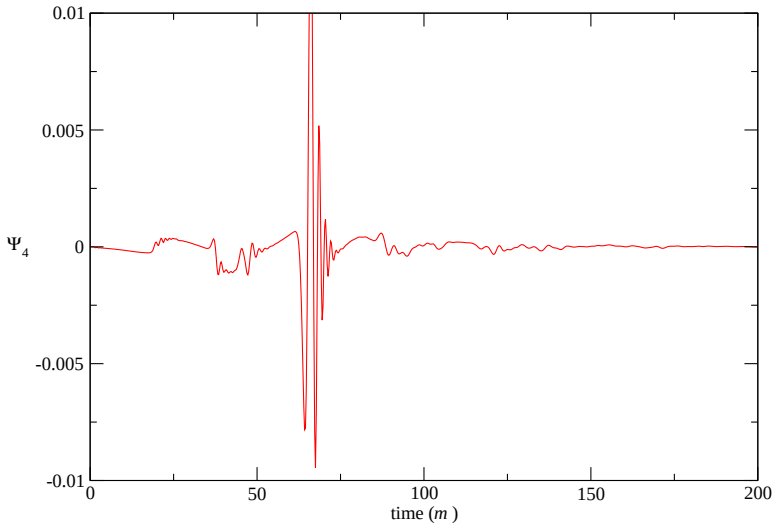

Figure 5.10: Ψ_4 from $t = 0m$ to $200m$. I have truncated a sharp peak of 0.016 for clarity.

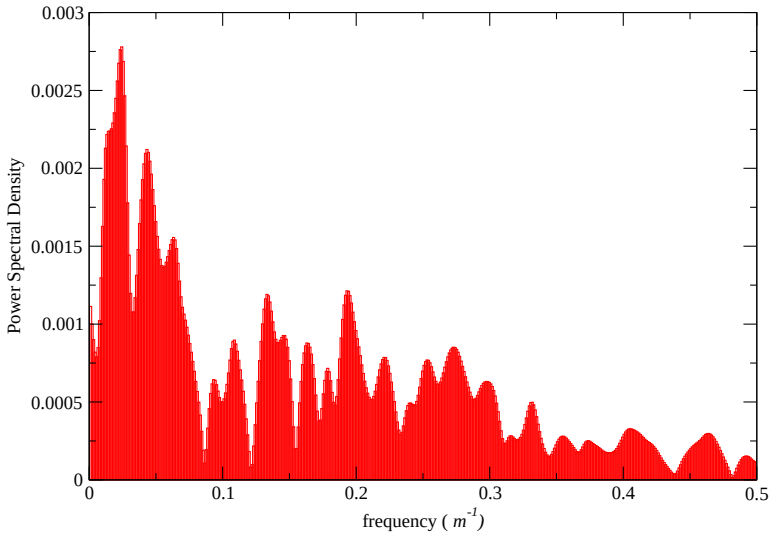

Figure 5.11: Discrete Fourier Transform of Ψ_4 from Fig 5.10 from $t = 100m$ to $1000m$.

the domain, so that the result is pure ring-down. Even so, the spectrum is much richer than that observed for the small-separation runs (compare with Figs 4.7 and 4.9). There is still a strong central peak, with a central period value of $T = 40.95m$, implying a mass $M = 2.44m$. The second peak is at $T = 23.1m$, and the third is at $T = 15.81m$. Somewhere between these two lies the $\ell = 4$ mode for the mass above, at $18.9m$.

$T_{obs}(m)$	(n, ℓ)	$M_T(m)$	% error	$M_\tau(m)$	% error
40.82	(1,2)	2.43	—	2.45	-0.8
22.99	(1,4)	2.96	21.8	1.41	-41.9
15.87	(1,6)	3.06	26.1	1.43	-41.2
10.70	(1,8)	2.73	12.6	2.78	14.2
9.22	(1,10)	2.92	20.2	2.21	-9.1
7.50	(1,12)	2.85	17.4	2.02	-17.1

Table 5.1: Mode fits to DFT from Fig 5.11. In each case, the peak position of the (1,2) mode is assumed correct, and errors from other modes and their widths are calculated relative to this.

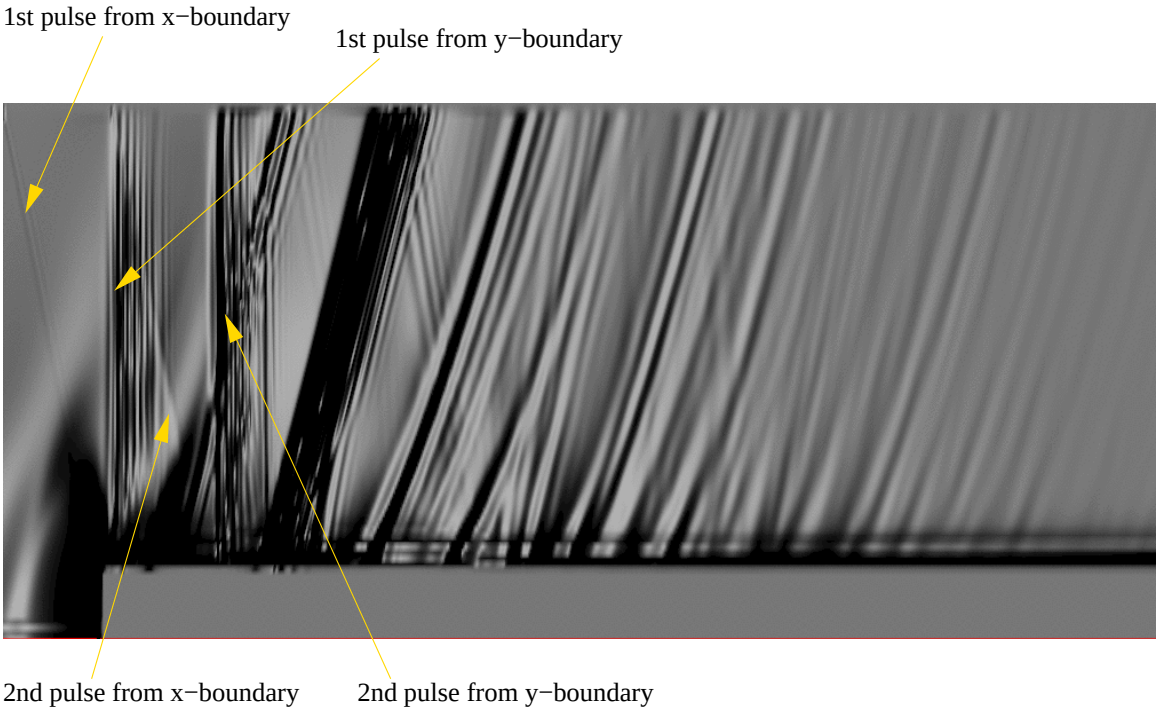

Figure 5.12: $t - x$ surface plot of Ψ_4 over the first $100m$ of evolution. Time is to the right, while the orthogonal direction depicts the x-axis.

5.5.4 Effect of Boundaries and Slicing Changes

Fig 5.12 shows a space-time surface similar to Fig 4.11 for the small-separation run. The outer boundary pulse effect from the small-separation runs is clearly visible here also; now, however, it occurs twice. The first occurrence is at the beginning of the evolution, when radiative outer boundary conditions are used for the BSSN evolution variables. The second pulse kicks in at the moment when the slicing fields α (or rather, Q) and β^i go live and also switch to radiative conditions.

That this can occur for slicing fields at all – quantities with no physical significance – is perhaps a bit troubling. On the other hand, a change in α and β^i causes a change the local coordinate scheme, and a sudden jump in time may affect the observed wave forms for a detector in a Doppler-shift-like action.

5.5.5 Alternative Outer Boundary Conditions

In Section 4.3.5, I showed that for the small-separation runs, varying the background IEF mass had no effect on the initial pulse traveling in from the outer boundaries at the start of the evolution. In the head-on case, for the most part, we use a background of the unsolved binary data, since we are very far from the small-separation case. We may ask if this choice

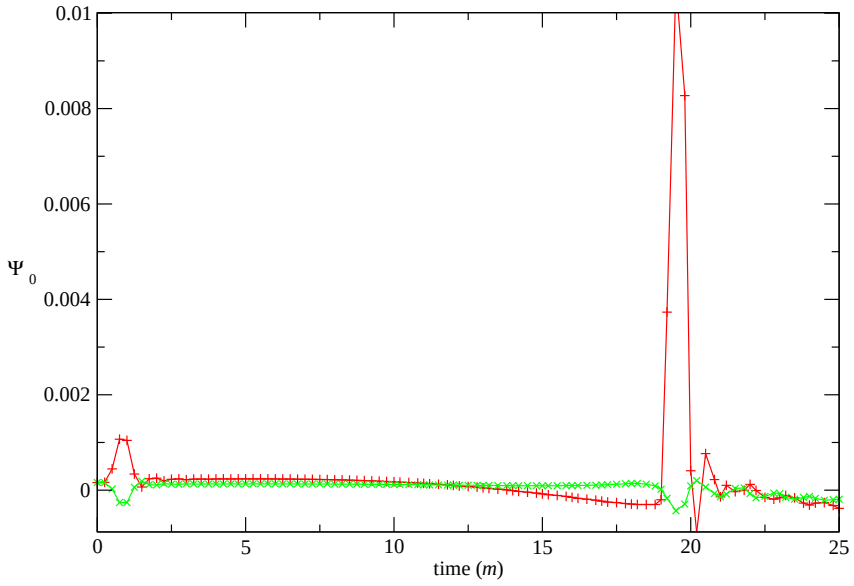

Figure 5.13: Early Ψ_0 wave forms for radiative outer boundary conditions on IEF (red +) and HuMaSh (green $\times$).

is important. As a brief comparison, I show in Fig 5.13 a plot of Ψ_0 at a fixed position ($x = 17.0m$) near the outer boundary.

It is interesting to note that using a HuMaSh background makes a definite difference – the initial pulse is only around half the size of that for an IEF background, and the later aggregated pulse (at $t = 19m$) is much better behaved. Coupled with the null result of Section 4.3.5, this leads us to believe that an important factor in the size of the pulse is the match of the background to the dipole shape of the binary data.

5.5.6 Alternative Slicing Merger Choices

Keeping the initial outer boundary condition fixed, I now look at the choice of when to merge the singularities for the slicing conditions. I mentioned in Section 5.4.2 that the effective meaningful window for switching to the live gauge conditions extends from the point of mask merger to the point of trajectory intersection. With this in mind, I performed two runs, identical up until the point of mask merger. At this point, one run switches to live gauge immediately, while the other waits until the merger of the trajectory, a few m later. These two choices can be contrasted by looking at a suitable quantity near the outer boundary.

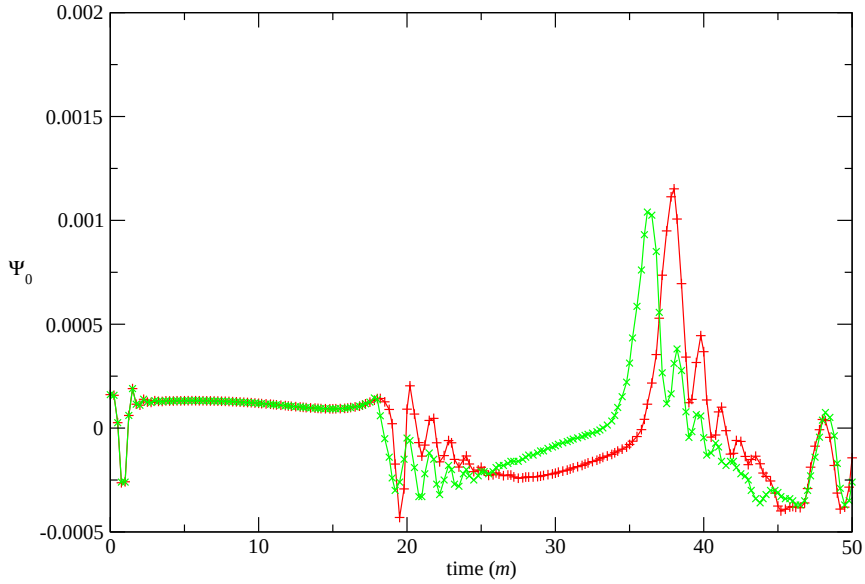

Figure 5.14: Early Ψ_0 wave forms for radiative outer boundary conditions for early slicing switch (green $\times$) and the late switch (red $+$).

Fig 5.14 shows Ψ_0 for both cases.

5.5.7 Very-Large-Separation Mergers

In many ways, the most impressive demonstration of our current abilities is that of the holes with the largest initial separation, $32m$. The run parameters can be found in the Appendix, Section C.1.

For this data, the holes are initially centered at points $(0, 0, \pm 16m)$. Of course, the coordinate positions do not give an invariant measure of distance. A better way of thinking of this separation is that in Kerr-Schild-type data, where the holes have individual horizons at radius $R_h = 2m$, the individual horizons are a distance $28m \approx 14R_h$ apart to begin with. This is comparable to the most widely separated puncture-data runs I am aware of [18].

Despite the use of octant symmetry, we are restricted in the domain size available; given the wide hole separation, the domain was chosen to be longer in the z-direction ($24m$) than in the transverse directions ($16m$). As can be seen in Section C.1, we used a larger boost ($v_N = \mp 0.25$) than is justified by Newtonian in-fall. This was necessary to ensure that the holes reached merger in a timely manner: for a lower initial velocity, the difference between

the gaussian-fit mask center and any trajectory of the form (5.4) grows larger than the excision mask radius. At this point, the singularities calculated in Q and β^i will fall outside the excision region, and lead to an immediate code crash.

The first common horizon was found at $t = 44m$, with an effective mass $M_{AH} = 1.97m$, and asymmetry $\epsilon = -30\%$. By the time mask merger took place, at $t = 47m$, $M_{AH} = 2.04m$, and the horizon had circularized substantially: $\epsilon = -6.7\%$.

Fig 5.15 shows the profile of K along the x- and z-axes at four times during this evolution: the initial time slice ($t = 0m$), just before merger ($t = 47.35m$), just after merger ($t = 49.0m$), and late in the ring-down stage ($t = 200m$). The difference between the third and fourth profiles indicates to what extent K has settled down in apparently isotropic coordinates.

This achievement is impressive not only because we have merged such widely separated holes, but also because single boosted-hole evolutions with the Maya code have died after about $50m$ of evolution time, at least, for significant boosts ($v = 0.5$).

5.5.8 The Advantage of Current Tracking Methods

I would like to present here some evidence of how well the combination of excision mask tracking and pseudo-Newtonian singularity trajectories performs for binary interactions. The differences show up best in cases of large boosts, where the asymmetry in K across the mask is the greatest.

The runs I have presented here do not have particularly large boosts initially. The most promising example is the $32m$-separation run above. The standard run was with a pseudo-Newtonian singularity trajectory and GaussTracker-centered masks, omitting the on-boundary mask points in the fit. As comparison cases, I re-ran the simulations (a) using the on-boundary points for the GaussTracker fits, and (b) ignoring GaussTracker, and centering using the pseudo-Newtonian trajectory, as is done for the singularities.

To indicate how differences occur between these cases, I plot part of the pre-merger z-trajectories of the mask center in Fig 5.16. The trajectories are very close for the bulk of the in-fall, but the boundary-less method consistently re-centers a little earlier than the with-boundary case, while the pseudo-Newtonian trajectory, which was fitted to the early with-boundary mask motion, deviates from both Gaussian-centered paths at late times. Between the two Gaussian paths, there is no clear winner until the "hiccough" at $t = 47m$. Fig 5.17 shows the profile of K along the z-axis at the time $t = 47.35m$, just before merger. Note that (a) the no-boundary case's mask is closer to the origin – it will merge at the next time-step; (b) the no-boundary case is both less peaked and smoother than the with-boundary case. In comparison, the Newtonian trajectory case isn't moving fast enough to contain the singularity, and in fact, that run crashed before it could complete merger.

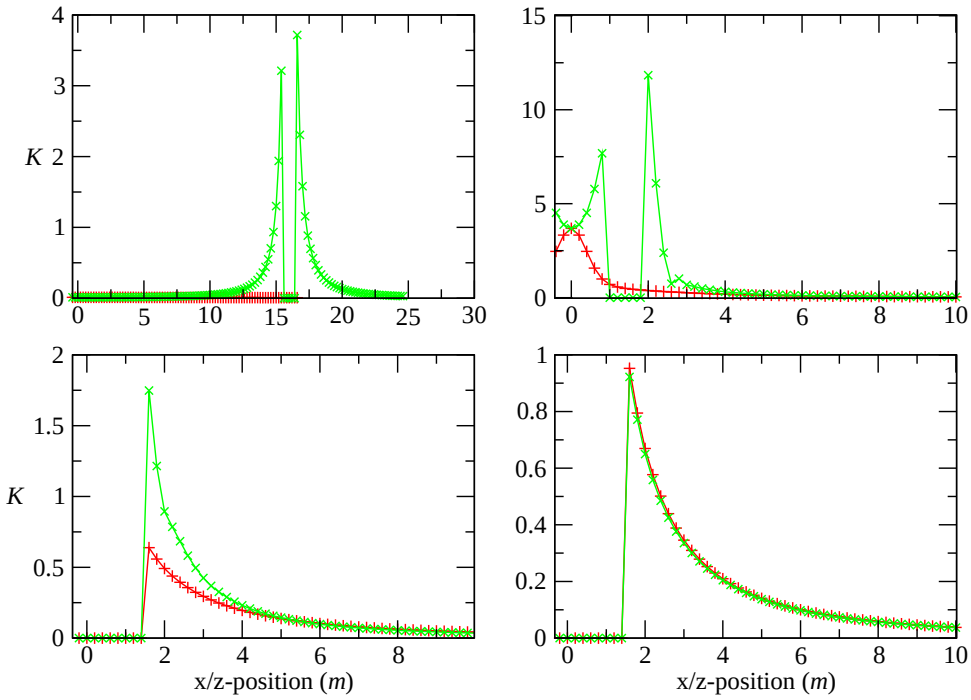

Figure 5.15: K at several times during the far-limit evolution: the initial time slice, just before mask merger, just after mask merger, and at late times. The data is shown for both the x-axis (red +) and the z-axis (green $\times$) at all times except the initial one.

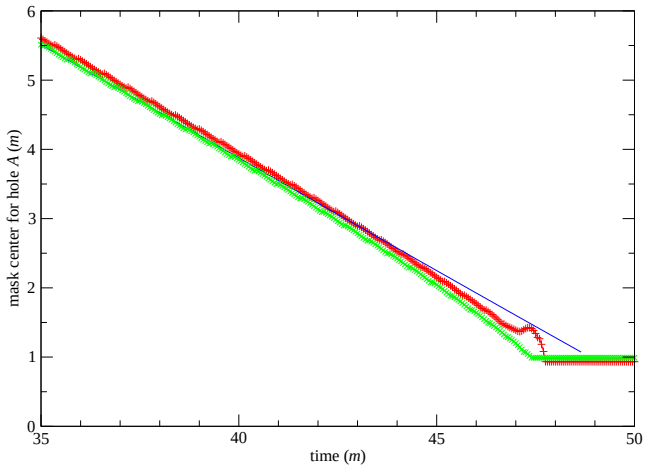

Figure 5.16: Position of hole A's mask center on the z-axis before merger. The three trajectories are Gaussian tracking using boundaries (red +), Gaussian tracking without boundaries (green $\times$), and Newtonian centering (blue).

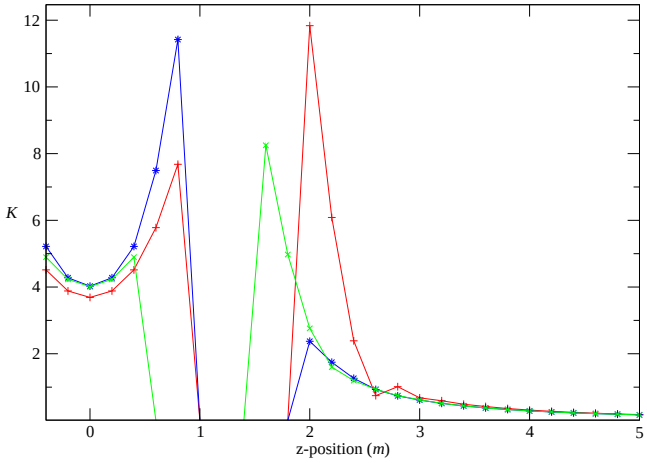

Figure 5.17: K along the z-axis at time $t = 47.35m$, shortly before merger. The three profiles are Gaussian tracking using boundaries (red +), Gaussian tracking without boundaries (green $\times$), and Newtonian centering (blue $*$).

5.6 Discussion

In this chapter, I've presented head-on merger evolutions, over a range of initial data sets, using superposed Kerr-Schild black-hole initial data.

The primary result from these evolutions is the success of the runs themselves: initially well-separated holes move toward each other across the computational grid under a combination of initial boosts and mutual attraction; they develop a common apparent horizon when sufficiently close; they ring down to a stationary single black hole; and this hole evolves stably for thousands of m. To my knowledge, these are the first long-term stable merger simulations of black holes with non-conformal initial data and singularities moving across the grid.

As well as achieving stability, I have extracted some basic information about the end-state black hole, and the gravitational radiation produced during the highly non-linear interaction. At field points near the outer boundary of the domain, wave forms are clearly visible, and of the right qualitative shape. The use of several different diagnostics – the apparent horizon effective mass, the monopole Weyl scalar Ψ_2 profile, the ADM energy, and the quasi-normal mode frequencies – has allowed several different measures of the final black hole mass. The variability of these estimates is indicative of the reliability of these diagnostics.

The success of these simulations is of interest for other reasons, however: as mentioned in the introduction, almost all full non-linear binary evolutions to date have been carried out using conformal data sets with the Bowen-York prescription for adding linear and angular momentum to the system. Studies carried out in the last three decades have generated gravitational wave forms from conformal data in the perturbative regime, in 2D axisymmetry and in full 3D, with good agreement in extracted wave forms. Conformal data, however, was often used for its convenience rather than its physical realism. If the wave forms produced by conformal data evolutions are physically correct and generic to black-hole mergers, then they should be visible in evolutions of different types of binary black hole data.

The comparative robustness of these evolutions leads us to believe that we may be able to address more physically relevant problems in the near future. The simulations presented here have a few obvious limitations, most of which should be resolved before the head-on collisions can be relied on to produce reliable wave forms. I will briefly address some of these below.

5.6.1 Areas for Improvement

First, and most directly, we are severely restricted by the size of the physical domains we can model. The MAYA code used for these investigations is simple and extensible, but is

to date, still a unigrid, Cartesian code. This keeps us firmly in the strong-field region of the interactions, meaning that our attempts to use effective radiative boundary conditions will never succeed properly; unphysical outer boundary reflections will always re-enter the BBH interaction region. Closely related to this is our inability to extract useful wave forms – pollution is unavoidable.

The most direct resolution of this problem is the introduction of some kind of mesh refinement into the code. Just four levels of refinement will allow the outer boundary to extend to around $r = 100m$, already an immense improvement. It is probable that we will soon be able to use the mesh refinement capabilities of the new Carpet thorns in CACTUS. With mesh refinement, we can both increase the initial separation of the black holes, and push our outer boundaries to much greater distances, where radiative boundary conditions are more appropriate.

Mesh refinement may also aid in the complementary area of mask merger: even assuming that our resolution is currently good enough to resolve all the necessary classical physics, the minimal mask-mask separation is determined not by physical length scales, but by the need for a fixed number of interpolation points between the masks.

A more natural merger procedure may have to be developed; one that would allow for non-regular excision regions. A post-merger excision region that could include the pre-merger masks in a minimal manner (like making a pill shape from two spheres) would be ideal, but involves more infrastructure than is present in the MAYA code at the moment.

Another area of ambiguity lies in our coordinate assumptions post-merger. Although the initial data used is constructed from the simple superposition of two Kerr-Schild black holes, we cannot guarantee that our late-time post-merger coordinates are Kerr-Schild coordinates of the final black hole. We can be confident that we are still using horizon-penetrating coordinates (as an apparent horizon can be found). We can also be sure from the symmetries of this particular problem that the end-state is symmetric about the numerical z axis. This still leaves ambiguity in the polar-radial mixing, and in the overall scaling of radial distances. As noted toward the end of Section 4.4, such coordinate warping could be confused with real physical constraint violations. It would also limit the usefulness of the Ψ_2 profile fits for the mass of the hole.

It is unclear how exactly to deal with this coordinate warping. If one can locate an apparent horizon at late times, the deviations from sphericity of the horizon shape could be used as a measure of the warping. An extension of this kind of idea has been suggested recently by Erik Schnetter [112], using surfaces of constant expansion to foliate the numerical space-time and analyze the gauge effects.

Considerable work has also been done in this area as part of the "Lazarus" project [17], where full 3D binary evolutions of limited duration are mapped at late times onto a perturbative evolution of a single end-state black hole. The accuracy of this mapping

depends strongly on a thorough understanding of the coordinates of the 3D system just before mapping.

Finally, and critically for more complex interactions than the simple head-on, we need a dynamic slicing condition that moves the lapse and shift in a sensible manner up to merger. Ideally, such a prescription would need no direct specification of the singularity positions and boosts. This, coupled with an outer boundary sufficiently far away that the metric is only a perturbation of flat space there, would remove any need for pre-knowledge of the hole parameters.

5.6.2 The Future

As well as being an independent check on past conformal data results, the Kerr-Schild head-on collisions here are a natural precursor to more complex and interesting binary interactions. In particular, the simulation of almost-circular orbits and slow binary inspiral to merger has been the Holy Grail of numerical relativity for several years. Most success in this area to date – again, with conformal data – has involved the use of co-moving coordinates with a quasi-stationary ansatz [36]. Even with co-moving coordinates, it seems likely that to achieve several orbits of the binary before plunge, dynamic singularity excision may be needed[3]. Our own recent progress with moving singularities through the computational domain [123] and the mergers presented in this thesis suggests that the conceptually more direct approach of orbiting singularities moving through the grid might be a real possibility in the near future.

[3]Private communication with B. Brügmann.

Chapter 6

Black Hole Spectroscopy

The material from this chapter is drawn from work done with Olaf Dreyer, Lee S. Finn, David Garrison, Ramon Lopez-Aleman and Badri Krishnan, and was published as [58].

This work was supported through the Center for Gravitational Wave Physics, which is funded by the NSF under cooperative agreement PHY 01-14375, NSF awards PHY 00-90091 and PHY 00-99559, the Eberly research funds of Penn State and the Albert-Einstein Institut.

6.1 Introduction

The numerical relativist's goal of black hole evolutions is not merely an academic exercise. There is a growing community of experimentalists involved in the search of direct evident of strong-field gravitational effects. Chief among these efforts is the international collaboration to observe gravitational waves produced by interactions of massive bodies (solar mass, massive and supermassive black holes, neutron stars). Ground-based detectors are already up and running, taking scientific data in urgent need of analysis. Space-based detectors may well be launched in the next decade or so. All these detectors require strong guidance from theorists and numericists about what to expect.

The formation of a black hole is the ultimate manifestation of strong field gravity. Although we lack detailed information about the gravitational radiation produced through most of the formation process, our knowledge of the near-end point affords us important knowledge about the nature of general relativity.

During the late stages in the aspherical formation of an astrophysical black hole the gravitational waves emitted are dominated by a set of quasi-normal modes (QNMs) [49, 50, 51]: exponentially damped sinusoids, whose frequency and damping times are characteristic of the black hole's mass and angular momentum[1]. Should we observe a QNM from a black hole and know also which particular normal mode we are observing we can determine, from the mode's frequency and damping time, the black hole mass and angular momentum [60, 66].

[1] And electric charge, as well; however, astrophysical black holes, which are our interest here, have negligible charge-to-mass ratio.

If, on the other hand, we observe several different QNMs from the same source and find that they are inconsistent with the spectrum predicted by general relativity in the sense that they *cannot* be explained by a single value of the mass and angular momentum we may infer that we are not observing a black hole. Here we develop this observation into a experimental test of the existence of black holes or, alternatively, general relativity.

Future observations by the Laser Interferometer Space Antenna (LISA) [55] offer us a different perspective on this question. LISA is expected to observe mergers of compact objects with masses in the range $10^6 - 10^8\,M_\odot$ [68]. In our present understanding, these compact objects can only be black holes. Observations by LISA of QNMs inconsistent with black holes would also be a test of the general relativistic no-hair theorem since an inconsistency in this mass range with black hole sources would indicate that physical scales other than mass and angular momentum were involved in the generation of the radiation.

In either sense the test described here is of general relativity based on gravitational wave observations. Eardley et al [59] proposed the first test of general relativity using gravitational wave observations. They investigated the polarization modes of gravitational waves in various metric theories of gravity and described how to identify the modes experimentally and use those observations to identify the spin content of dynamical gravity. The first actual test of general relativity relying on its prediction of the existence of gravitational waves was made by Taylor & Weisberg [126]. They described how the observed orbit and orbit decay of the Hulse-Taylor binary pulsar PSR B1913+16 led to a strong consistency check on the predictions of general relativity. Finn [65] proposed a different test of the spin content of dynamical gravity, based on the possibility of a space-based detector in circumsolar orbit observing the induction-zone field associated with solar oscillations. Ryan [108] has outlined how observations of the gravitational radiation from capture orbits of solar mass compact bodies about a supermassive black hole may allow the determination of certain multipole moments of the central hole, thereby testing the prediction of general relativity. More recently, Will [136] and Finn & Sutton [67, 125] have described tests of general relativity that bound the mass of the graviton, and Scharre & Will [110] and Fairhurst et al [62] have shown how gravitational wave observations of pulsars may be used to bound the value of the Brans-Dicke coupling constant.

The preceding tests can be grouped into three different classes. One set of tests, which includes [67, 108, 110, 126], is based on energy conservation arguments: the observed evolution of a system or of the radiation from a system, is related to the energy loss expected owing to the radiation. A second class of tests, which includes [59, 62, 65], focuses on the observed polarization modes of the field. A third class, which includes [136], involves the frequency-dependent dispersion relationship associated with a massive graviton. The test described here is of a new class, based on the unique character of the radiation spectrum associated with a disturbed black hole.

Damped sinusoidal motion is ubiquitous for systems approaching equilibrium and one expects that collapse or coalescence will lead, in any theory of gravity, to some form of QNM ringing. If we observe a QNM spectrum that is inconsistent with an isolated black hole, then there are two possibilities. On the one hand, general relativity may yet be correct, but we are not observing an isolated black hole approaching equilibrium[2]. Similarly, we may be observing the radiation arising from a compact body that is not a black hole — e.g., a neutron star [8], a boson star [114] or strange-matter star [4], whose QNM spectrum will be determined by the properties and configuration of the appropriate matter fields — or a black hole carrying a previously unknown macroscopic charge e.g., a dilaton field [74]. On the other hand, general relativity may not be the correct theory of gravity in the strong field limit. Thus, while no single observation may rule out general relativity, a set of observations, each of a different source, *none* of which is consistent with an isolated black hole, could suggest the need to consider alternative theories of gravity in the strong-field limit.

This chapter is organized as follows. In Section 6.2 we briefly describe the QNMs of a Kerr black hole and explain how the idealized observation of two or more modes in the absence of noise enables us to extract the mass and angular momentum of the black hole. Real boats, of course, rock, and Section 6.3 generalizes the discussion to include experimental errors and describes how one can use noisy gravitational wave observations of QNMs to test general relativity. In Section 6.4 we demonstrate, in a model numerical simulation, the use of this method as applied to LISA observations. Section 6.5 investigates the range to which we can expect LISA to observe sources strong enough for this test to be applied. We conclude in Section 6.6 with a summary of our main results.

6.2 Ideal Observations

6.2.1 Quasi-normal Modes of Kerr Black Holes

Following the aspherical collapse to a black hole, one expects that the final black hole can be described as the perturbation of a stationary Kerr hole. The dominant part of the gravitational waves emitted as the black hole settles down can be described as a sum over a countably infinite set of damped sinusoids, each characterized by an amplitude, phase, frequency and damping time. (At still later times, the radiation will be dominated by power-law tails arising from the backscatter of radiation off the space-time curvature in the neighborhood of the black hole [49, 50, 51]; however, here we are interested in the earlier, and higher amplitude, QNM ringing.) In this sub-section we review those properties of the black

[2]In fact, the uniqueness theorems have only been proved for vacuum space-times and they are not true in the presence of arbitrary matter fields or radiation. Nevertheless, it would be a great surprise if the space-time in the vicinity of a black hole is not close to Kerr in some approximate sense.

hole QNM spectrum that are important for our investigation; more detailed examinations of the spectrum itself can be found in [21, 22, 88, 90, 99, 100].

QNMs appear as solutions to the equations describing perturbations of a stationary black hole space-time, subject to the boundary conditions of no in-going radiation from infinity and no up-coming radiation from the horizon. The perturbation equations describing Schwarzschild black holes were first described by Regge & Wheeler [107] and Zerilli [143, 144]. The first QNM solutions to these equations were found by Vishveshwara [[134]]. Teukolsky found the corresponding perturbation equations for Kerr black holes [127, 128], and, with Press, first investigated their QNM solutions [104].

Gravitational wave detectors respond to a linear combination of the radiation in the two polarization modes of the incident gravitational waves. The observable strain $h(t)$ in, for example, the arms of an interferometric detector may be written, for QNMs, in the form

$$h(t) = \mathrm{Re}\left[\sum_{n,\ell,m} A_{n\ell m} e^{-i(\omega_{n\ell m} t + \phi_{n\ell m})}\right] \tag{6.1}$$

where the summation indices characterize the particular mode, which is related to the angular dependence of the mode amplitude and phase on a sphere of constant (Boyer-Lindquist) radius about the black hole through ℓ and m, and the "harmonic" through the index n: $\ell = 2, 3 \ldots$, $|m| \leq \ell$ and $n = 1, 2 \ldots$. For the Schwarzschild geometry the symmetry is spherical, the appropriate decomposition of the metric perturbation is given by the usual spherical harmonics, and modes differing only in m are degenerate. For Kerr the symmetry is axisymmetric and the orthonormal decomposition of the perturbation is by spheroidal harmonics [128]. The amplitudes $A_{n\ell m}$ and phases $\phi_{n\ell m}$ depend on the initial conditions and the relative orientation of the detector and the source; however, the complex frequency $\omega_{n\ell m}$ depends only on the intrinsic parameters of the underlying black hole: i.e., its mass M and angular momentum aM^2. (We assume that the black hole carries no significant electric charge.)

For fixed a the complex frequency $\omega_{n\ell m}$ scales as M^{-1}; thus, we define the dimensionless frequency $\Omega_{n\ell m}$,

$$\Omega_{n\ell m} \quad := \quad M\omega_{n\ell m} := \left(2\pi F_{n\ell m} + \frac{i}{T_{n\ell m}}\right) \tag{6.2}$$

where $F_{n\ell m}$ and $T_{n\ell m}$ are the real dimensionless frequency and damping time of the modes. The corresponding physical frequency $f_{n\ell m}$ and damping time $\tau_{n\ell m}$ are given by

$$\omega_{n\ell m} = 2\pi f_{n\ell m} + \frac{i}{\tau_{n\ell m}} = 2\pi \frac{F_{n\ell m}}{M} + \frac{i}{M T_{n\ell m}}. \tag{6.3}$$

(We use geometrical units with $G = 1$ and $c = 1$.) The dimensionless $\Omega_{n\ell m}$ (or $F_{n\ell m}$ and $T_{n\ell m}$) depend *only* on the also dimensionless black hole angular momentum parameter a. Fig 6.1 shows $\Omega_{n\ell m}$ as a function of a for $n = 1, 2$ and $\ell = 2, 3$, and $|m| \leq \ell$.

6.2.2 From Quasi-normal Modes to Testing Relativity

If we observe only one mode, characterized by its complex frequency ω (cf. eq. 6.3), what can we say about the underlying black hole?

Corresponding to the observed ω is the line $\Omega = M\omega$, $M \in \mathbb{R}_{\geq 0}$, in the dimensionless Ω plane (cf. eq. 6.2). Such a line is shown in Fig 6.1. This line will intersect some subset of the family of $\Omega_{n\ell m}$ curves, characteristic of black hole normal modes. Each intersection corresponds to a black hole mass M, angular momentum parameter a, and mode $n\ell m$ consistent with the observed ω. Knowing only f and τ, then, we cannot uniquely identify the black hole mass and angular momentum, but we can reduce the possibilities to a (possibly countably infinite) set of (a, M) pairs. If we knew $n\ell m$ as well, we would know a and M exactly.

Now suppose that we observe two modes from the same black hole, each characterized by its own frequency and damping time. Fig 6.2a shows, in schematic form, the line $M\omega$ for each of the two modes (denoted $+$ and $\times$) and their intersection with several different $\Omega_{n\ell m}$ curves in the complex Ω plane. Corresponding to each mode is a set of candidate (a, M) pairs that may describe the underlying black hole. Each candidate mass and angular momentum parameter is a point in the (a, M) plane, as shown in Fig 6.2b. With two or more modes, there must be at least one common candidate mass and angular momentum.

This is, in essence, our proposal for a test: interpreting the observation of several normal modes ω_k, $k \geq 2$, as arriving from a single, general relativistic black hole, we conclude that general relativity is self-consistent if the observed ω_k are consistent with at least one black hole (a, M). If no such (a, M) exists for the observed ω_k either we have observed something other than an isolated black hole or we have a contradiction with the predictions of the theory.

(As an aside, it is possible (though unlikely) that we get more than one value of (a, M) consistent with the observed frequencies. This can happen if we have two mode pairs $(n_1, \ell_1, m_1; n_2, \ell_2, m_2)$ and $(\tilde{n}_1, \tilde{\ell}_1, \tilde{m}_1; \tilde{n}_2, \tilde{\ell}_2, \tilde{m}_2)$ which give rise to the same frequency ω. In this case the observations would still be consistent with general relativity though we could not use that observation to measure M and a. The important point of our test is the existence of at least one (a, M) pair consistent with the observations.)

Noise and other experimental realities ensure that there will be no *exact* agreement between the observed ω_k and a general relativistic black hole even if general relativity is correct. The challenge, then, in developing a practical test is to determine when the differences between the candidate (a, M) pairs associated with the different observed modes

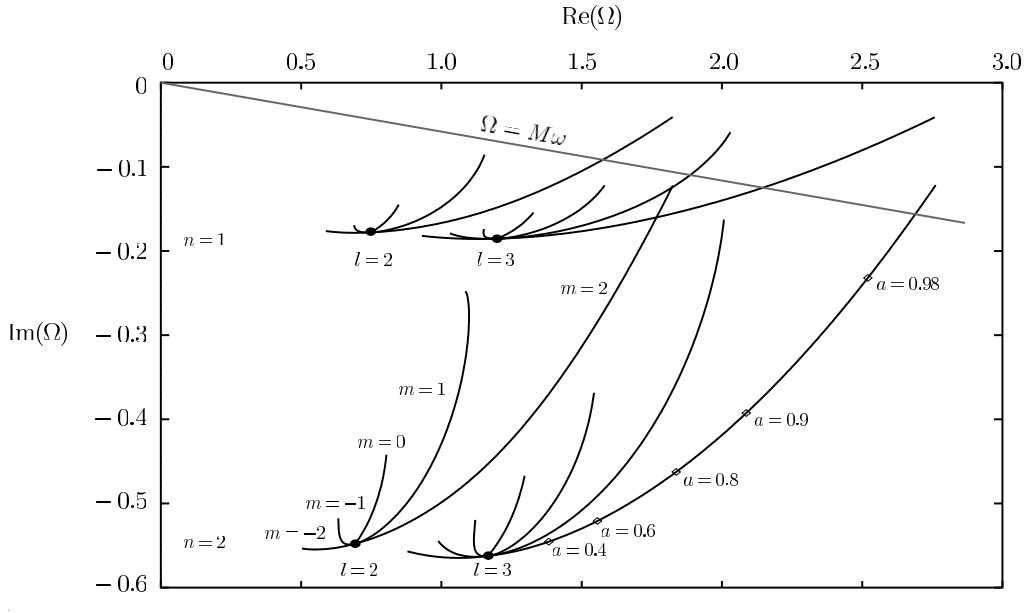

Figure 6.1: The dimensionless, complex QNM frequencies $\Omega_{n\ell m}$ for rotating, uncharged black holes. Each family of curves corresponds to one $n\ell$ pair, and each branch to a possible value of m. The large black dot at the base of each family is the Schwarzschild ($a = 0$) limit, where the frequencies are degenerate in m. This degeneracy is broken for $a \neq 0$, and the curves emanating from the dots give the QNM frequencies for Kerr black holes as a function of positive a for different m. In this figure a ranges from 0 to 0.9958, with the small diamonds marking the QNM frequencies for $a = 0.4, 0.6, 0.8, 0.9$, and 0.98. In this figure, an observation, corresponding to a (complex) frequency ω, is represented by the line $\Omega = M\omega$, parameterized by the (unknown) black hole mass M. Each intersection of this line with a QNM curve in dimensionless Ω represents a candidate $n\ell m$, M and a for the mode.

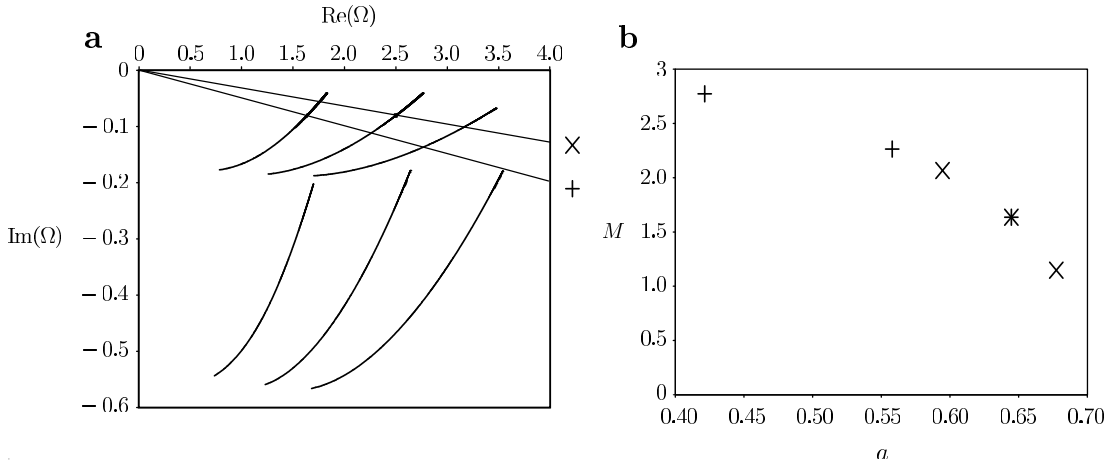

Figure 6.2: Here we show, in schematic form, several $\Omega_{n\ell m}(a)$ curves and their intersection with the lines $M\omega_i$, $M > 0, i = 1, 2$, corresponding to two observed modes. We denote these two lines by $+$ and $\times$, respectively. (**b**) The candidate (a, M) pairs determined in figure (a) are plotted here in the (a, M)–plane. The pairs belonging to ω_1 are denoted by $+$, the ones belonging to ω_2 by $\times$. There is only one candidate (a, M) consistent with both observations — indicated by the overlapping $+$ and $\times$ — and this is the actual mass and angular momentum of the underlying black hole.

are so great as to be statistically inconsistent with general relativity. In the next section we face this challenge.

6.3 A Test of Relativity

6.3.1 A Reformulation of the Test

Before we discuss the role that noise plays in our analysis it is helpful to reformulate the test described in section 6.2.2 and Fig 6.2. Consider an ordered N-tuple of QNMs,

$$\mathcal{Q} := \{(n_k, \ell_k, m_k) : k = 1 \ldots N\}\,. \tag{6.4}$$

Each $\mathcal{Q}$ may be regarded as a function that maps M and a to a set of observable frequencies

$$\mathcal{Q}(a, M) := \left\{M^{-1}\Omega_{n_k\ell_k m_k}(a) : k = 1 \ldots N\right\}. \tag{6.5}$$

Each N-tuple $\mathcal{Q}$ thus describes a two dimensional surface in the $(2N+2)$-dimensional space $\mathcal{S}$,

$$\mathcal{S} := (a, M, \omega_1, \ldots, \omega_N), \tag{6.6}$$

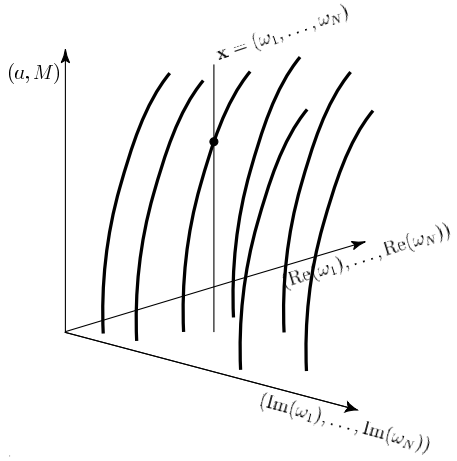

Figure 6.3: A reformulation of the consistency criterion. A set of quasi-normal modes $\mathcal{Q} = \{(n_k, l_k, m_k) : k = 1, \ldots, N\}$ corresponds to a surface in the $(2N + 2)$-dimensional space depicted in this figure. A measurement $\omega = (\omega_1, \ldots, \omega_N)$ is consistent with general relativity if the constant surface that is obtained by ranging over all (a, M) while keeping the frequencies ω fixed intersects at least one of the surfaces corresponding to one of the sets $\mathcal{Q}$. This intersection is indicated in this figure by a dot.

with different N-tuples corresponding to different sets of N modes. (In section 6.3.3 we will understand the ω_k to represent observed QNM frequencies and damping times.)

An observation ω consists of an N-tuple

$$\omega := (\omega_1, \ldots, \omega_N)\,. \tag{6.7}$$

The observation ω also corresponds to a surface in $\mathcal{S}$. The observation is consistent with a black hole if the surface of constant ω intersects one of the surfaces $\mathcal{Q}$. Fig 6.3 shows a low-dimensional projection of such an observation ω together with several surfaces (which appear as curves) for different N-tuples $\mathcal{Q}$. A moment's consideration should convince one that this new criterion is equivalent to the criterion formulated above in section 6.2.2.

In practice the situation is less ideal: noise distorts our observation, so that — even if we are observing black hole QNMs — the measured ω will not intersect a curve $\mathcal{Q}$. In the remainder of this section we describe how our ideal test is made practical and meaningful for real observations.

6.3.2 Confidence Intervals and Testing General Relativity

In a frequentist analysis, the observation, the sampling distribution, an ordering principle and a probability combine to determine a confidence interval. In this section we use this construction to form a confidence region in the (a, M)-plane, given a noisy observation ω.

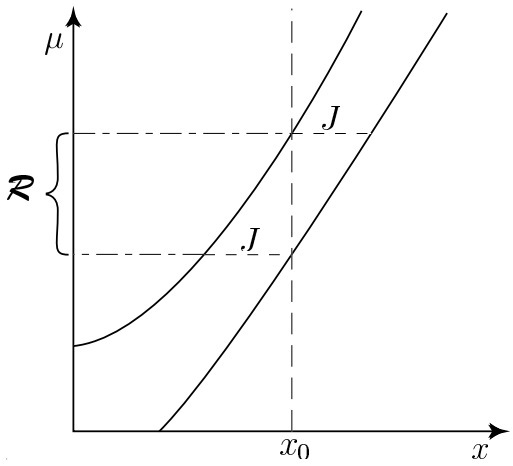

Figure 6.4: The construction of classical confidence intervals. A sampling distribution $P(x|\mu)$, an ordering principle R, and a probability p are needed to construct a confidence interval. The ordering principle is used to find the intervals $J(\mu)$ such that $\int_{J(\mu)} dx\ P(x|\mu) = p$. The classical confidence interval $\mathcal{R}$ is then given by the set of μ for which $J(\mu)$ contains the measured value x_0.

We begin by reviewing the construction of a classical confidence interval for the one-dimensional case following [98] (alternatively, see e.g.[124]). We suppose that we make measurements of a random variable x from which a quantity μ is determined. The sampling distribution $P(x|\mu)$ is the probability of making the observation x given a particular μ. Formally, an *ordering principle* is a function $R(x|\mu)$, which we use to identify a sub-interval J of x according to

$$J(\mu|r) := \{x : R(x|\mu) > r)\}\,. \tag{6.8}$$

The parameter r is chosen such that the region $J(\mu|r)$ encloses a fixed probability p:

$$\int_{J(\mu|r)} P(x|\mu)\,dx = p\ . \tag{6.9}$$

Given an observation x_0, the probability-p *confidence interval* is the range of μ for which $J(\mu|r(p))$ includes x_0 as shown in Fig 6.4. In an actual experiment, the choice of the value of the parameter p is made by the experimentalist. Typical choices are 90%, 95% or 99%.

The choice of the ordering principle $R(x|\mu)$ is a key ingredient in the construction of confidence intervals. Different choices will lead to different confidence intervals for the same observation: for example, one choice of ordering principle will always determine intervals of the form $(-\infty, x)$, while another choice will always determine intervals of the form (x, ∞). Neither choice is *a priori* right or wrong. Here we will choose $R(x|\mu) = P(x|\mu)$ so that the intervals are given by level surfaces of the distribution $P(x|\mu)$. The main advantage of this

ordering principle is that it is simple and it works in any dimension. Consider, for example, a two-component observation depending on one parameter a. There is, as before, a sampling distribution $P(x, y|a)$ and an ordering principle $R(x, y|a) = P(x, y|a)$. Confidence intervals can be defined in the same way as in the case of a one-dimensional observation; the interval J is now a two-dimensional region. Since this system is over-determined — we are now trying to determine *one* parameter a by measuring *two* quantities x and y — the measured x and y will have to satisfy additional constraints in order to give a non-vanishing confidence region. This is in fact precisely what happens in the black hole quasi-normal mode problem: any single measurement of ω can be explained by *some* (a, M), but a measurement of two or more ω can be simultaneously consistent with at least one (a, M) pair only if the no-hair theorem is true and the modes arise from a single black hole.

We can now describe our test of relativity. Note that not all observations ω will lead to a non-empty confidence interval: i.e., for some ω there will be no (a, M) consistent with the observation. We frame our test in terms of the confidence interval we construct: if we make an observation ω for which there is no probability p confidence interval, then we say that the observed normal modes are inconsistent with an isolated black hole with confidence p. Conversely, if there does exist a probability p confidence interval, then we have verified that general relativity is self-consistent at this confidence level.

Finally, we should point out an aesthetic flaw of our choice of ordering principle. The function $P(x|\mu)$ is a density and, therefore, not invariant under a reparameterization of x. If we were to use a new parameter $x' = f(x)$ for some smooth monotonic function f, the confidence region obtained for μ using a measurement of x may not coincide with the region obtained using a measurement of x'. In the one-dimensional case, there exists another ordering principle based on the likelihood ratio [63] which is reparameterization invariant; however, we have not been able to generalize this to higher dimensions. While aesthetically displeasing, there is nothing wrong with the choice we have made, which is natural given the physical association of the parameters M and a with the black hole mass and angular momentum.

6.3.3 Generalization to Quasi-normal Modes

The generalization to QNM observations is straightforward. Each observation consists of N complex QNM frequencies ω_k and associated amplitude signal-to-noise ratios ρ_k, which characterize both the amplitude of the signal at that frequency and the uncertainty in the determination of ω_k (cf. [66]):

$$\omega := (\omega_1, \ldots, \omega_N) \tag{6.10}$$

$$\rho := (\rho_1, \ldots, \rho_N)\,. \tag{6.11}$$

For definiteness suppose that ω_k and ρ_k are identified via maximum likelihood techniques [66]. There is a minimum signal-to-noise associated with each mode, which is set by the requirement that the observation must identify N modes.

Observations ω corresponding to a black hole characterized by (a, M) and signal-to-noise ρ_k are distributed according to the sampling distribution

$$P(\omega|a, M, \mathcal{Q}, \rho) := \left(\begin{array}{l} \text{Probability of making observation} \\ \omega \text{ given the actual } N\text{-tuple } \mathcal{Q} \\ \text{and signals-to-noise } \rho. \end{array} \right). \tag{6.12}$$

In general the sampling distribution depends upon the nature of the detector noise and the analysis procedure that identifies the modes ω_k. For large signal-to-noise ratios it will generally reduce to a multivariate Gaussian in $\mathrm{Re}(\omega_k)$ and $\mathrm{Im}(\omega_k)$ and for smaller signal-to-noise ratios it can be determined via simulation.

Now consider the region of the space $\mathcal{S}$ (cf. section 6.3.1) defined by

$$P(\omega|a, M, \mathcal{Q}, \rho) > p_0 \tag{6.13}$$

with p_0 such that

$$p = \int_{P(\omega|a,M,\mathcal{Q},\rho)>p_0} P(\omega|a, M, \mathcal{Q}, \rho) d^{2N}\omega \tag{6.14}$$

for a fixed p. We say that the observation ω is consistent with a black hole if the actual observation ω is included in this region for some (a, M). Fig 6.5 illustrates the comparison of an observation with the region defined by equations (6.13–6.14). It remains only to specify p.

To help in specifying p it is useful to examine more closely its meaning. Suppose we have chosen a value of p. That value of p determines a confidence region. Now consider an ensemble of identical detectors, each observing simultaneously the same black hole event and its corresponding QNMs. The fraction of these observations that does not intersect the confidence region is the *false alarm probability* $\alpha(p)$, so-called because it is the probability that an observation will be falsely deemed to be inconsistent with a black hole. The probability α is a monotonic function of p; therefore, we can specify α in lieu of p. For observations whose characteristic frequency corresponds to masses greater than neutron star masses, which we are confident originate with black holes, we propose setting p so that $\alpha(p)$ — the probability of falsely rejecting the hypothesis that we have in fact observed a black hole — is small (e.g., less than 1%). In other words, the standard of evidence for declaring that we have discovered "new physics" should be high.

The false alarm probability function $\alpha(p)$ will depend on the signal strength, as characterized by the signal-to-noise ratios; consequently, it will need to be determined on an

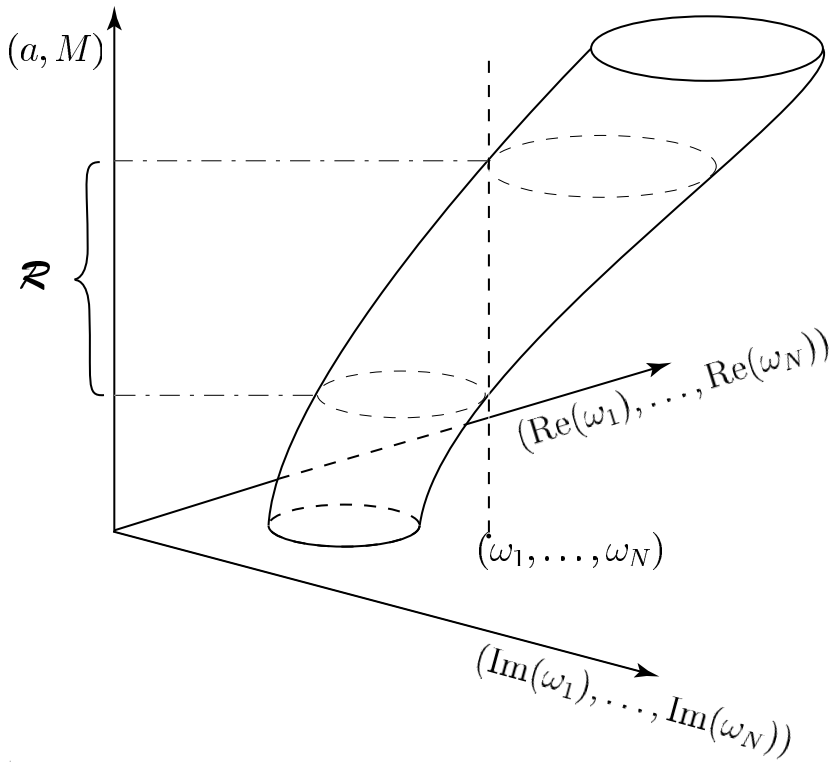

Figure 6.5: The construction of classical confidence intervals generalized to higher dimensions. Given a sampling distribution P, an ordering principle, and a probability p one can construct classical confidence regions R just as in the one-dimensional case. The difference here is that we are now trying to determine a small number of parameters (a, M) from a larger number of observations $\omega = (\omega_1, \ldots, \omega_N)$. There are thus additional consistency conditions that need to be satisfied to obtain a non-empty confidence region R. This is the basis of our proposed test.

observation-by-observation basis. Thus the calculation of $\alpha(p)$ by a Monte Carlo simulation is the final ingredient we need. In the next section we demonstrate the test through a numerical example where we calculate $\alpha(p)$.

6.4 A Numerical Example

In the previous section we described a general procedure for testing general relativity by observing QNMs. In this section we explore its effectiveness numerically through a set of simulated observations drawn from a hypothetical black hole population population inspired by potential LISA observations, and a hypothetical population of non-black hole compact object sources, or NBHs. (We say "inspired" because, in fact, for the purpose of this analysis the observations are characterized entirely by the dimensionless signal-to-noise ratio and

mode quality factor, with the dimensioned mode frequency simply setting a scale. Thus, the conclusions we reach are as valid for LISA observations as they are for observations at the same signal-to-noise with ground-based detectors.)

For the BH observations we find the relationship between the false alarm probability α and the probability p that appears in equation (6.14). For the NBH observations there are no "false alarms": every observation is of something not a black hole. Instead, there are *false dismissals*: observations that we mistakenly classify as consistent with a black hole. The probability of a false dismissal, denoted β, depends on the choice of p or, alternatively, the choice of false alarm probability $\alpha(p)$ that we make for the purpose of defining the test. (The false dismissal probability depends also on how the spectrum of BHs and NBHs differ.) The smaller the false dismissal probability the more sensitive the test is to discovering "new physics" or identifying non-black hole sources. For the NBH observations we evaluate the false dismissal probability as a function of the false alarm probability.

6.4.1 Mode Detection

Our concern here is with the question of statistical inference from QNM observations, characterized by their signal-to-noise, frequency and damping time. We do not venture to explore how, from a signal-processing standpoint, these events are identified and so characterized and none of the qualitative conclusions drawn in this section depend on the method of event identification and characterization. Nevertheless, it is worthwhile to comment briefly on the challenges associated with detecting and characterizing damped sinusoidal signals.

Matched filtering is often invoked as the preferred method for detecting signals in noise when the signal is known exactly, up to a few parameters. This does not mean, however, that matched filtering is the most computationally efficient means of detection. Similarly, while matched filtering may be optimal for these problems, it does not mean that there is necessarily a large difference in efficiency between matched filtering and other less optimal methods of event identification. Matched filtering stands out only when the signal being sought has significant defining features that allow it to be discriminated from noise events. This is not the case for damped sinusoids with low quality factor $\mathcal{Q}$.

In fact, matched filtering is known to be a very poor way to search for damped sinusoids. A matched filter search for a damped sinusoid of unknown frequency and damping time in a time series $h(t)$ is equivalent to taking the Laplace transform of $h(t)$, which is computationally difficult to do accurately. (A set of 'matched filter templates', consisting of damped sinusoids, applied to a time series $h(t)$ corresponds to a sample of the Laplace transform of h.) Fortunately, the problem of analyzing noisy data for damped sinusoidal signals is not unique to gravitational-wave physics, and other methods, beyond matched filtering, exist (cf, e.g., [80, 115]).

While none of the results below depend on the method of detection, we chose, where

choices need to be made, to use theoretical results from matched filtering studies to describe the uncertainties associated with frequency and damping time determination. Such choices do not affect our qualitative conclusions and, in any event, relevant quantitative conclusions could not be drawn without both LISA data and the choice of a specific analysis method.

6.4.2 Simulating Black-Hole QNM Observations

For definiteness we focus on observations of two QNMs. For the purpose of illustration we consider black hole masses and angular momenta consistent with potential observations by the LISA detector [68]. We first draw an (a, M) pair from the distribution

$$P(a, M) = P(a)P(M) \tag{6.15}$$

$$P(a) \propto \begin{cases} 1 & \text{for } a \in [0, 0.986) \\ 0 & \text{otherwise} \end{cases} \tag{6.16}$$

$$P(M) \propto \begin{cases} M^{-1} & \text{for } M \in (2.5 \times 10^5 M_\odot, 4.5 \times 10^8 M_\odot) \\ 0 & \text{otherwise} \end{cases} \tag{6.17}$$

The range of M is determined by the frequency band where LISA is expected to be most sensitive; the range of a is determined by the maximum angular momentum expected of a black hole spun-up by thin-disk accretion [133].

Corresponding to each (a, M) pair we choose the QNMs corresponding to $(n = 1, \ell = 2, m = 2)$ and $(n = 1, \ell = 4, m = 4)$. We assign each mode the same signal-to-noise ratio, which we treat here as sufficiently large that the errors associated with the measurements are normally distributed with covariance matrix C_{ij} equal to the inverse of the *Fisher information matrix* I_{ij} (see e.g. [124]) as given in [66, equation 4.14]. This is in fact a mathematical lower bound — the *Cramer-Rao bound* — on the covariance matrix. We draw from this error distribution errors in the frequencies and damping times that we add to the "real" frequencies and damping times to determine the simulated observations: noisy QNM frequencies and damping times.

Given this pair of QNM frequencies and damping times with errors we ask whether the two modes are in fact observationally distinguishable: if the frequencies and damping times are not sufficiently different then no real observation would ever result in the given pair. For instance, the five $(n = 1, l = 2)$ modes are degenerate at $a = 0$; consequently, no matter how large the signal-to-noise ratio, if a is sufficiently small it is impossible to resolve these five modes observationally.

To decide whether the two modes we are investigating are observationally distinguishable we invoke a "resolvability criterion": denoting the frequencies (damping times) of the two

modes as f_1, f_2 (τ_1, τ_2) we say that the two modes are distinguishable if

$$|f_1 - f_2| > \frac{1}{\min(\tau_1, \tau_2)} \,. \tag{6.18}$$

We discard any mode pair that does not satisfy this criterion.

The result is an *observation,* which consists of a pair of signal-to-noise ratios and associated frequencies and damping times. (The observation does *not* include knowledge of black hole mass or angular momentum, or the $n\ell m$ associated with the frequencies or damping times.)

6.4.3 False Alarm Probability α

For each simulated observation ω, constructed as described in section 6.4.2 we evaluate the smallest probability $p = p_{\rm min}$ such that equations (6.13) and (6.14) describe a region $\mathcal{S}$ that covers ω for some (a, M). The false dismissal fraction $\alpha(p)$ is the fraction of $p_{\rm min}$ determinations that are greater than p: i.e., the fraction of BH observations that we would reject as originating from a black hole for threshold p.

Ideally, in evaluating p we would consider every possible $n\ell m$ for each ω_k. In practice, we consider only a finite subset of low-order (in both n and ℓ) modes, corresponding to our expectation that these are the modes most likely to be excited to large amplitude. In our simulations we considered only modes corresponding to $(n = 1, \ell = 2, m = 0)$, $(n = 1, \ell = 2, m = 2)$, $(n = 1, \ell = 3, m = 3)$ and $(n = 1, \ell = 4, m = 4)$. Since for these simulations we observed two distinguishable QNMs there were twelve possible ordered pairs of modes. Fig 6.6 shows α as a function of p for four different signal-to-noise ratios. Each $\alpha(p)$ curve is constructed from 10^4 simulated observations with that amplitude-squared signal-to-noise in each mode.

6.4.4 False Dismissal Probability Calculation

Complementary to α, the probability that we incorrectly decide we have observed QNMs from something other than a black hole, is the probability that we falsely conclude we have observed QNMs from a black hole. This probability is referred to as the false dismissal probability and commonly denoted β.

The false dismissal probability depends on the detailed character of the source, which is not a black hole. Strong gravitational wave sources are compact, with radius R on order their mass GM/c^2 and oscillations periods of order GM/c. At the frequencies where LISA will have its greatest sensitivity — 10^{-2}–10^{-4} Hz, corresponding to masses of order 10^6–10^8 $M_\odot$ — we know of no compact sources that are not black holes. For the purpose of illustration and to give a sense of the ability of the test described here to "discover" new

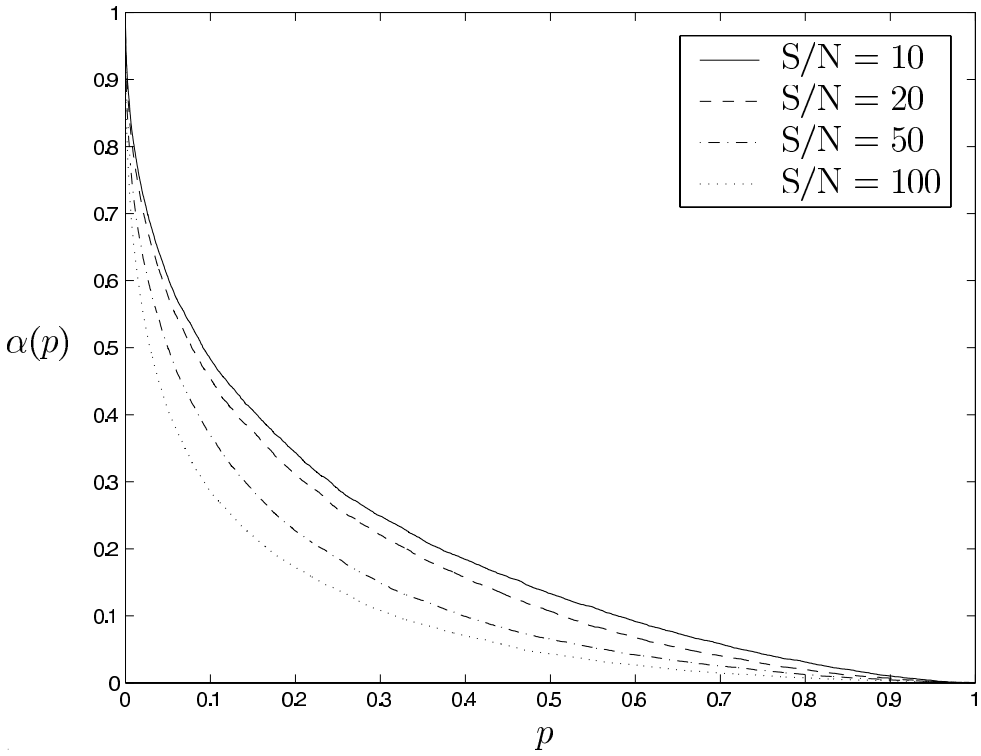

Figure 6.6: False alarm probability α as a function of the probability p appearing in equation (6.14). A false alarm is a misidentification of a QNM pair as arising from something other than a general relativistic black hole.

physics, we suppose a population of sources whose frequencies and damping times share the same relationship as certain neutron star w-modes calculated in [8]. Referring to [8, table 1, col. 1, lines 3, 5] we consider observations consisting of two modes

$$\begin{aligned} M\omega_1 &= 0.471 + 0.056i, \\ M\omega_2 &= 0.654 + 0.164i, \end{aligned} \tag{6.19}$$

where M is drawn from the distribution given in equation (6.17). In exactly the same way that we used simulations in section 6.4.3 to determine α as a function of p we calculate from these simulations β as a function of p. Together $\alpha(p)$ and $\beta(p)$ determine $\beta(\alpha)$, which we show in Fig 6.7. A measure of the effectiveness of the test is the degree to which the curves for different signal-to-noise fall below the $\beta = 1 - \alpha$ diagonal. (A "test" that randomly picked a fraction α of observations as not black holes would have $\beta = 1 - \alpha$. Any "test" that can do better than randomly choosing in this way will have a $\beta(\alpha)$ curve that falls below this diagonal.) As expected the test also does better with stronger signals. Consider a false alarm threshold of 1%. Then for observations with $\rho^2 = 10$ we have a better-than-40% chance of distinguishing NBH sources from BH sources. This climbs to better-than-90% chance for observations with $\rho^2 = 100$.

6.5 Potential for Application

We have shown that, given at least two QNM signals, from the same source and with sufficiently large signal-to-noise, we can cleanly distinguish black holes from other astrophysical sources. We have also demonstrated how this can be used as a test of general relativity. In this section we investigate the potential for application of this test in future LISA observations by asking, first,

- How distant can LISA-scale black-hole sources be and still have multiple QNMs detected at sufficiently high signal-to-noise?

and, secondly,

- What is the rate of sources that we may expect within this distance?

We focus attention on each individual QNM. The signal strength, characterized by the signal-to-noise ratio at the detector, depends on the energy radiated in the mode, the radiation pattern associated with the mode, and the relative orientation of the detector and the source. Following [69, eq. 2.30] we can average over these latter angles to obtain the mean-square signal-to-noise associated with the $n\ell m$ mode as a function of the mode energy

$$\langle \rho^2 \rangle = \frac{2(1+z)^2}{5\pi^2 D(z)^2} \int_0^\infty df \frac{1}{f^2 S_n(f)} \frac{dE_e}{df_e}[(1+z)f] \tag{6.20}$$

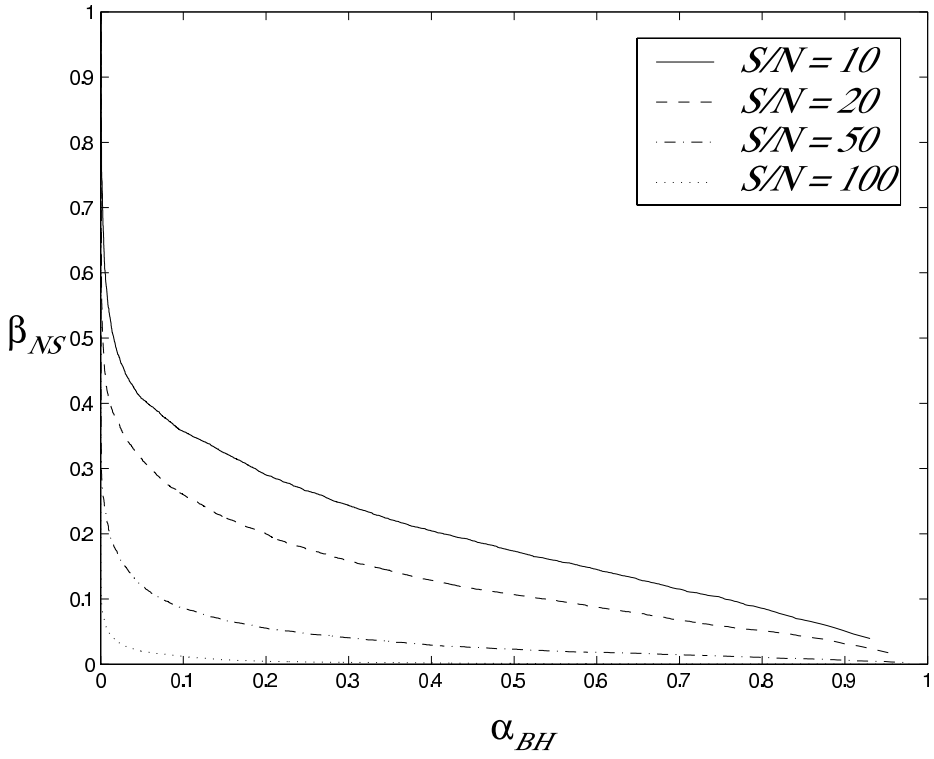

Figure 6.7: False dismissal probability as a function of false alarm probability $\beta(\alpha)$. The false dismissal probability depends on the non-black hole QNM spectrum, which we have taken to have the same ratio of frequencies and relationship between frequencies and damping times as neutron star w-modes.

where z and $D(z)$ are, respectively, the redshift and the luminosity distance to the source.

For the mode, we assume the form

$$h_{n\ell m}(t) = A_{n\ell m} \exp\left(-\frac{\pi f_{n\ell m} t}{Q_{n\ell m}}\right) \sin\left(2\pi f_{n\ell m} t\right) \tag{6.21}$$

where $Q_{n\ell m} \equiv \pi f_{n\ell m} \tau_{n\ell m}$. Note that $Q_{n\ell m}$, which is an observable property of a QNM, is independent of source redshift, while the observed $f_{n\ell m}$ and $\tau_{n\ell m}$ depend on redshift.

The ringdown energy spectrum of the $n\ell m$ mode is taken from eq (3.18) of [69]:

$$\begin{aligned}\frac{dE_e}{df_e} &= \frac{\epsilon_{n\ell m}}{F_{n\ell m}} \frac{Q_{n\ell m}}{\left(4Q_{n\ell m}^2 + 1\right)} \frac{M^2 f^2}{\pi^3 \tau^2} \left[\frac{1}{[(f - f_{n\ell m})^2 + (2\pi\tau)^{-2}]^2}\right. \\ &\quad \left. + \frac{1}{[(f + f_{n\ell m})^2 + (2\pi\tau)^{-2}]^2}\right],\end{aligned} \tag{6.22}$$

where the mode amplitude $A_{n\ell m}$ has been replaced with the fraction $\epsilon_{n\ell m}$ of the mass radiated in that mode, defined by

$$\epsilon_{n\ell m} := \frac{1}{M} \int_0^\infty \frac{dE}{df} df. \tag{6.23}$$

Using this spectrum in the formula above, and approximating the LISA noise power spectral density $S_n(f)$ as constant over the signal band we integrate over frequencies and invert the result to obtain an approximate range over which we can observe a mode $n\ell m$ with signal-to-noise greater than $\rho^2_{n\ell m}$:

$$D(z)^2 = \frac{8}{5\pi^2} \frac{Q_{n\ell m}^2}{4Q_{n\ell m}^2 + 1} \frac{(1+z)^3 M^3}{F_{n\ell m}^2} \frac{\epsilon_{n\ell m}}{S_n \rho^2_{n\ell m}} \frac{G^3}{c^7}\,. \tag{6.24}$$

The relationship between luminosity distance and redshift we take to be given by eqs (23), (25) of [41]. This relationship depends on cosmological parameters that characterize the universe and its expansion and for these we use the values determined by the first season's WMAP observations [122]. Thus, given a threshold $\rho^2_{n\ell m}$, black holes radiating a fraction $\epsilon_{n\ell m}$ of their rest energy in mode $n\ell m$ are observable with a redshift z satisfying eq (6.24).

Numerical simulations suggest that energy emitted in QNMs during ringdown may be of order 1% of the rest-mass energy of the hole [32, 85]. For equal-mass black hole mergers, the simulations suggest that the $\ell = 2$ modes will be by far the strongest, with total emitted energies greater than the $\ell = 4$ modes by as much as three orders of magnitude (see [11]). Thus we may assume that the weaker mode of a QNM pair carries away a fraction 10^{-5} of the final black-hole mass. To be more conservative, we consider instead a considerably smaller emitted mass fraction $\epsilon_{n\ell m} = 10^{-7}$.

LISA will be most sensitive in the frequency band $10^{-3} \leq f \leq 10^{-2}$ Hz, where the noise power spectral density is expected to be 5×10^{-45} Hz^{-1}. Black holes at redshift z whose (low-order) QNM frequencies peak in this band have a mass of order $\mathcal{M}/(1+z)$, where $\mathcal{M}$ ranges from 10^6 to $10^7 M_\odot$. Focusing just on these black holes we find, from equation (6.24), that LISA can expect to see pairs of QNMs associated with black-hole mergers, with the weaker mode having a signal-to-noise $\rho > 10$, within a redshift of ~ 52 (for extremal-spin Kerr) or ~ 36 (for Schwarzschild); that is, LISA will be able to observe mergers associated with the assembly of essentially all galaxies throughout the universe. Considering a broader frequency band of $10^{-1} - 10^{-4}$ Hz will increase the mass range (and thus the rate of observed mergers) further with the addition of more and less massive (red-shifted) black-hole mergers, though to a somewhat smaller redshift.

To summarize, we expect that pairs of QNMs associated with the mergers of supermassive black holes in the red-shifted mass range of $10^6 - 10^7 M_\odot$ will be observable by LISA with significant signal-to-noise throughout the observable universe as long as the weaker mode carries away at least a fraction 10^{-7} of the final black-hole mass. The rate of such black-hole mergers depends on redshift owing to evolution and the scaling of the intrinsic black-hole mass corresponding to observed QNM frequency with redshift. There is considerable uncertainty at present about the event rate for such observations; however, estimates out to $z = 20$ – far smaller than the range to which LISA is sensitive – range from 0.3 to 100 per year [52, 77, 78]. Correspondingly, we expect that this test will find application in forthcoming LISA observations.

6.6 Discussion

We have described a qualitatively new test of the existence of general relativistic black holes, based on the gravitational radiation they emit when they are formed or when they are impulsively excited, e.g., through a merger event. Radiation from an impulsively excited black hole, such as might arise in the course of a non-spherical black hole formation event or the coalescence of a black hole with another black hole or compact object, has a component that consists of a sum of damped sinusoids. This signature is, characteristic of the radiation from any impulsively excited, damped source. For any given mode, the scale of the frequency and damping time measures the black hole mass and angular momentum. Similarly, the relationship of the different modes to each other — i.e., the spectrum — is unique to black holes. We have described here how this relationship can be used to test the proposition that observed gravitational waves, characteristic of an impulsively excited, damped source, in fact originate from a general relativistic black hole. Such a test can be characterized in at least two different ways: as a definitive "proof" that a black hole has been observed, or as a test of the so-called "no-hair" theorem of general relativity.

To demonstrate the effectiveness of this test we have evaluated numerically the probability that the test will mistakenly fail to identify an actual black hole. By introducing a hypothetical gravitational wave source whose characteristic frequencies and damping times are similar to those of neutron star w-mode [8] we have also evaluated numerically the probability that the test will incorrectly identify w-mode oscillations of a neutron star, or any object whose spectrum is similar to a black hole. Together these results demonstrate that for sources with the signal-to-noise expected of, for example, massive black hole coalescences detected by LISA, the test proposed here can cleanly discriminate black hole sources. Finally, we have shown that LISA can be expected to observe signals of this kind and strength throughout the universe with a rate that may be as great as 100/y.

This method can be used to measure mass and angular momentum of a black hole. Using gravitational waves to measure mass and angular momentum is an idea that has been around for some time [60, 66]. In these previous works it was assumed that the mode observed was of a known order (e.g., the mode with the longest damping time, or the lowest order, etc.). With the observation of two or more modes the requirement that a single mass and angular momentum explain the complete set likely permits the mass and angular momentum to be determined uniquely.

The field of gravitational-wave detection is new. The current generation of ground- and space-based gravitational-wave detectors is opening a new frontier of physics: *gravitational-wave phenomenology,* or the use of gravitational wave observations to learn about the physics of gravitational-wave sources and gravity itself. We are only just beginning learn how to exploit the opportunities it is creating for us. As gravitational-wave observations mature, we can expect more and greater recognition of their utility as probes of the character of relativistic gravity. The opening of this new frontier promises to be an exciting and revealing one for the physics of gravity.

Chapter 7

Conclusion

Much has changed in the few years since the last binary Kerr-Schild evolutions: we have an evolution system with better mathematical properties than the standard ADM system, which has made it possible to evolve a single black hole stably in 3D. We have a simple, but powerful, method for dealing with generic black hole singularities. With such progress, it was natural to re-address the head-on problem.

The results I have presented in this thesis show that we are able to bring two well-separated black holes to the point of merger, achieve that merger, and let the result of the merger ring down to a final quiescent state. We can do this – for the first time – in full 3D, without invoking all the symmetries available to us, and while using data that represents the initial and final black holes in an intuitive way. This is, in itself, cause for celebration.

More encouraging is the realization that since our methods – our singularity masking and tracking, the construction of our gauge and boundary conditions – have been generic, so the problems we can go on to address immediately are also generic. To evolve spinning black holes, or holes with an impact parameter, we need not rewrite our evolution code, or change our system of equations. Assuming satisfactory resolution of the outstanding problems mentioned in Section 5.6 – including dynamic slicing conditions and boundary effects – the natural next step will be grazing collisions of increasingly greater obliqueness. The obvious ultimate aim is to perform simulations of real binary orbits, with the individual holes moving across the numerical grid for several hundred – or thousand – m.

The procedures used here also highlight the continuing usefulness of the simplest kind of evolution code – the finite-difference mesh. At a time where many researchers are using spectral methods, finite elements or adaptive mesh refinement, we have been able to complete interesting simulations with a conceptually trivial unigrid code.

Given the progress we have made with these more modern methods, it seems likely that soon we will be able to carry out full evolutions for situations where the limited duration of 3D runs made it necessary to switch to lower-dimensional perturbative evolutions. At the very least, we can offer much longer 3D runs to give greater choice as to when to switch to the perturbative regime. Moreover, we will have an alternative, fully three-dimensional, path to the late-time behavior of the space-time, and will be able to supply consistency checks on the results of perturbative codes.

What I believe we are now seeing come to light is not just a new, marginally different way of solving an already-solved problem. I have presented evidence here that we have now perhaps reached the abilities previous generations of numerical relativists had thought were just around the corner. The relative simplicity of the tools used, and the robustness of the merger itself, give us hope that we are finally getting to the point where we can banish the demons that have beset us since we first attempted fully 3D evolutions. We still have much to occupy us in our search for the better identification of physical quantities, such as the mass and radiation spectrum of the final black hole; however, at least we can be quite sure that we will be able to get to that final black hole, and maintain stability for time periods much longer than is actually needed to extract physics.

Appendix A

Exact Black-Hole Solutions in the 3+1 Split

In this Appendix, I present some commonly used slicings for black hole space-times. The decompositions were performed by hand or – for the more complicated situations – using the Maple symbolic manipulation package and its built-in "Tensor" module. Much of the simplification had to be done by hand, however.

A.1 Spherical Symmetry – The Schwarzschild Solution

The Schwarzschild solution is the unique space-time describing spherical symmetry and asymptotic flatness in vacuum.

A.1.1 Schwarzschild Coordinates

The simplest set of coordinates used to describe spherically-symmetric space-time is that due to Schwarzschild himself [113]. Using spherical polar coordinates, the metric for a black hole of mass M is:

$$ds^2 = -\left(1 - \frac{2M}{r}\right) dt^2 + \left(1 - \frac{2M}{r}\right)^{-1} dr^2 + r^2\, d\Omega^2.$$

This yields the ADM fields:

$$\begin{aligned} \alpha &= \left(1 - \frac{2M}{r}\right)^{\frac{1}{2}}, \\ \beta^i &= 0, \\ \gamma_{ij} &= \mathrm{diag}\left(\frac{r}{r-2M}, r^2, r^2 \sin^\theta\right), \\ K_{ij} &= 0. \end{aligned} \tag{A.1}$$

In Cartesian coordinates, the four-metric becomes

$$ds^2 = -\left(1 - \frac{2M}{r}\right) dt^2 + \left(\delta_{ij} + \frac{2M}{r-2M}\frac{x^i}{r}\frac{x^j}{r}\right) dx^i dx^j,$$

yielding

$$\begin{aligned} \alpha &= \left(1 - \frac{2M}{r}\right)^{\frac{1}{2}}, \\ \beta^i &= 0, \\ \gamma_{ij} &= \delta_{ij} + \frac{2M}{r - 2M}\frac{x_i}{r}\frac{x_j}{r} \Rightarrow \gamma^{ij} = \delta^{ij} - \frac{2M}{r}\frac{x^i}{r}\frac{x^j}{r}, \\ K_{ij} &= 0, \end{aligned} \tag{A.2}$$

where $x_i \equiv \delta_{ij}x^j$ and $r \equiv \sqrt{\delta_{ij}x^i x^j}$.

A.1.2 Isotropic Schwarzschild Coordinates

The Schwarzschild solution in Isotropic Schwarzschild coordinates is (see [95], exercise 31.7):

$$ds^2 = -\left(\frac{2r - M}{2r + M}\right)^2 dt^2 + \left(1 + \frac{M}{2r}\right)^4 \left(dr^2 + r^2\, d\Omega^2\right),$$

or, in Cartesian coordinates,

$$ds^2 = -\left(\frac{2r - M}{2r + M}\right)^2 dt^2 + \left(1 + \frac{M}{2r}\right)^4 \delta_{ij} dx^i dx^j.$$

This yields the ADM fields

$$\begin{aligned} \alpha &= \frac{2r - M}{2r + M}, \\ \beta^i &= 0, \\ \gamma_{ij} &= \left(1 + \frac{M}{2r}\right)^4 \delta_{ij}, \\ K_{ij} &= 0. \end{aligned} \tag{A.3}$$

These coordinates are designed to emphasize the conformally-flat nature of the Schwarzschild solution - apart from an overall conformal factor, the three-metric is that of flat three-space.

In these coordinates, there is an inversion symmetry at the black hole horizon: we can switch to a new radial coordinate $r' \equiv M^2/4r$ and retain the form of the three-metric. Then as $r \to 0$, $r' \to \infty$, and we can view the interior of the hole not as a singularity but another asymptotically flat universe, causally disconnected from ours. The two asymptotically flat regions are connected by a "throat", narrowest at the horizon $r = r' = M/2$.

Numerically, this way of looking at the Schwarzschild solution is attractive, as by using an inversion-symmetric inner throat boundary condition, we can avoid dealing with the physical Schwarzschild singularity.

A.1.3 Ingoing Eddington-Finkelstein Coordinates

Most useful for our purposes are the horizon-penetrating Ingoing Eddington-Finkelstein coordinates [61, 64]. This is the simplest non-trivial instance of a Kerr-Schild-type metric; here $H = \frac{M}{r}$, and in spherical polars, $\ell_\mu = (1, 1, 0, 0)$, leading to the four-metric:

$$ds^2 = -\left(1 - \frac{2M}{r}\right) dt^2 + \frac{4M}{r} dt\, dr + \frac{r + 2M}{r} dr^2 + r^2 d\Omega^2.$$

The four-metric decomposes into:

$$\begin{aligned}
\alpha &= \sqrt{\frac{r}{r+2M}}, \\
\beta^i &= \left(\frac{2M}{r+2M}, 0, 0\right), \\
\gamma_{ij} &= \mathrm{diag}\left(\frac{r+2M}{r}, r^2, r^2 \sin^2\theta\right), \\
K_{ij} &= \mathrm{diag}\left(-\frac{2M(r+M)}{r^2\sqrt{r(r+2M)}}, 2M\sqrt{\frac{r}{r+2M}}, 2M\sqrt{\frac{r}{r+2M}}\sin^2\theta\right).
\end{aligned} \tag{A.4}$$

In cylindrical polar coordinates (q, ϕ, z), the four-metric becomes:

$$ds^2 = -dt^2 + dq^2 + q^2 d\phi^2 + dz^2 + \frac{2M}{r}\left(dt + \frac{q}{r}dq + \frac{z}{r}dz\right)^2,$$

while in Cartesian coordinates it becomes:

$$ds^2 = -\left(1 - \frac{2M}{r}\right) dt^2 + \frac{4Mx_i}{r^2} dt\, dx^i + \left(\delta_{ij} + \frac{2M}{r}\frac{x_i}{r}\frac{x_j}{r}\right) dx^i dx^j,$$

which leads to the standard ADM fields:

$$\begin{aligned}
\alpha &= \sqrt{\frac{r}{r+2M}}, \\
\beta^i &= \frac{2M}{r+2M}\frac{x^i}{r} \Rightarrow \beta_i = \frac{2Mx_i}{r^2}, \\
\gamma_{ij} &= \delta_{ij} + \frac{2M}{r}\frac{x_i x_j}{r^2} \Rightarrow \gamma^{ij} = \delta^{ij} - \frac{2M}{r+2M}\frac{x^i x^j}{r^2}, \\
K_{ij} &= \frac{2M}{\sqrt{r^3(r+2M)}}\left(\delta_{ij} - \frac{2r+M}{r}\frac{x_i x_j}{r^2}\right).
\end{aligned} \tag{A.5}$$

As this solution is so important for the work in the main body of the thesis, I have written down some additional useful quantities:

- the three-metric determinant: $|\gamma| = 1 + \frac{2M}{r}$,

- the three-Christoffel symbol: $\Gamma^i_{jk} = \frac{2Mx^i}{r^2(r+2M)}\left(\delta_{jk} - \frac{3x_jx_k}{2r^2}\right)$,
- the mixed-index extrinsic curvature: $K^i{}_j = \frac{2M}{\sqrt{r^3\,(r+2M)}}\left(\delta^i{}_j - \frac{2r+3M}{r+2M}\frac{x^ix_j}{r^2}\right)$,
- the three-Ricci tensor: $R_{ij} = \frac{M(r+4M)}{r^2(r+2M)^2}\delta_{ij} - \frac{M(3r+8M)x_ix_j}{r^4(r+2M)^2}$,
- the three-Ricci scalar: $R \equiv \gamma^{mn}R_{mn} = \frac{8M^2}{r^2(r+2M)^2}$,
- the trace-free extrinsic curvature: $A_{ij} \equiv K_{ij} - \frac{K}{3} = \frac{2M(2r+3M)}{3\sqrt{r^3(r+2M)^3}}\left[\delta_{ij} - \frac{3r+4M}{r}\frac{x_ix_j}{r^2}\right]$.

Finally, we can extract from the above the BSSN-decomposed fields:

$$\begin{aligned}
\phi &= \frac{1}{12}\log\left(\frac{r+2M}{r}\right),\\
\tilde{\gamma}_{ij} &\equiv e^{-4\phi}\gamma_{ij} = \left(1+\frac{2M}{r}\right)^{-1/3}\left[\delta_{ij} + \frac{2M}{r}\frac{x_ix_j}{r^2}\right],\\
K &\equiv K^m{}_m = \frac{2M(r+3M)}{\sqrt{r^3(r+2M)^3}},\\
\tilde{A}_{ij} &\equiv e^{-4\phi}A_{ij}\\
&= \left(1+\frac{2M}{r}\right)^{-1/3}\frac{2M(2r+3M)}{3\sqrt{r^3(r+2M)^3}}\left[\delta_{ij} - \frac{3r+4M}{r}\frac{x_ix_j}{r^2}\right],\\
\tilde{\Gamma}^i &\equiv -\tilde{\gamma}^{im}_{\ ,m} = \left(1+\frac{2M}{r}\right)^{1/3}\frac{8M(r+3M)x^i}{3r^2(r+2M)^2}.
\end{aligned} \tag{A.6}$$

A.2 Axial Symmetry - The Kerr Solution

After Schwarzschild, this is the simplest vacuum solution. It is parametrized by the black hole mass, M, and the *spin*, a, the angular momentum per unit mass (i.e., $a \equiv J/M$). When $a = 0$, the hole becomes spherically symmetric, with no angular momentum; that is, it becomes a Schwarzschild black hole in some coordinates. At the other end we have the *extremal* limit, where a attains its highest value, M (that is, the largest possible angular momentum is $J_{max} = M^2$).

A.2.1 Boyer-Lindquist Coordinates

In Boyer-Lindquist coordinates [29], the Kerr line element for a spin vector aligned along the z-axis looks like:

$$\begin{aligned} ds^2 &= -\left(1 - \frac{2Mr}{\Sigma}\right) dt^2 - \frac{4Mar\sin^2\theta}{\Sigma} dt\, d\phi + \frac{\Sigma}{\Delta} dr^2 + \Sigma d\theta^2 \\ &\quad + \sin^2\theta \left(r^2 + a^2 + \frac{2Ma^2 r \sin^2\theta}{\Sigma}\right) d\phi^2 \end{aligned}$$

Here $\Sigma \equiv r^2 + a^2\cos^2\theta$ and $\Delta \equiv r^2 - 2Mr + a^2$. This reduces to the Schwarzschild solution, *in Schwarzschild coordinates*, when $a = 0$.

A.2.2 Kerr-Schild Coordinates

More useful for our purposes is the *Kerr-Schild* form [84] of the Kerr solution. The four-metric is written:

$$g_{\mu\nu} = \eta_{\mu\nu} + 2H\ell_\mu\ell_\nu \Rightarrow g^{\mu\nu} = \eta^{\mu\nu} - 2H\ell^\mu\ell^\nu,$$

where ℓ_μ is a flat-space null vector: $\ell_\mu\ell^\mu \equiv \eta^{\mu\nu}\ell_\mu\ell_\nu = 0$. From this generic form, we can deduce something of the 3+1 decomposition:

$$\begin{aligned} \alpha &= \frac{1}{\sqrt{1 + 2H\ell_0^2}} \\ \beta_i &= 2H\ell_0\ell_i \ , \ \beta^i = \frac{2H\ell_0\ell^i}{1 + 2H\ell_0^2} \\ \gamma_{ij} &= \eta_{ij} + 2H\ell_i\ell_j \ , \ \gamma^{ij} = \eta^{ij} - \frac{2H}{1 + 2H\ell_0^2}\ell^i\ell^j \\ \Rightarrow \gamma_{ij,k} &= 2\left[H_{,k}\ell_i\ell_j + H\ell_{i,k}\ell_j + H\ell_i\ell_{j,k}\right] \\ K_{ij} &= \alpha\left[\ell_i H_{,j} + \ell_j H_{,i} + H\ell_{i,j} + H\ell_{j,i} + 2H^2\left(\ell_i\ell_m\ell_{j,m} + \ell_j\ell_m\ell_{i,m}\right)\right. \\ &\quad \left. + 2H\ell_i\ell_j\ell_m H_{,m}\right] \end{aligned} \tag{A.7}$$

In Kerr-Schild coordinates, a black hole with arbitrarily-directed spin vector has a four-

metric given by (see also [102]):

$$H = \frac{Mr^3}{r^4 + (\vec{a}\cdot\vec{x})^2}\ ,\ \ell_\mu = \left(1, \frac{r^2\vec{x} - r\vec{a}\times\vec{x} + (\vec{a}\cdot\vec{x})\,\vec{a}}{r\left(r^2+a^2\right)}\right) = (1, \vec{\ell})$$

and r is the solution of $r^4 - \left(\rho^2 - a^2\right)r^2 - (\vec{a}\cdot\vec{x})^2 = 0$ (where $\rho \equiv |\vec{x}|$):

$$r = \sqrt{\frac{\rho^2 - a^2 + \sqrt{(\rho^2-a^2)^2 + 4(\vec{a}\cdot\vec{x})^2}}{2}} = \sqrt{\frac{\rho^2 - a^2 + \Omega}{2}}$$

For zero spin, this reduces to the *Ingoing Eddington-Finkelstein* solution - the coordinates are horizon-penetrating. The most difficult part of the 3+1 split is the calculation of the extrinsic curvature. We do this in stages. First, the derivative of the "ring" distance w.r.t. the coordinate positions:

$$r_{,i} = \frac{1}{2r}\left(1 + \frac{\rho^2 - a^2}{\Omega}\right)x_i + \frac{\vec{a}\cdot\vec{x}}{\Omega r}a_i$$

Then the derivatives of the metric functions are

$$\begin{aligned} H_{,i} &= \frac{H}{r}\frac{(3(\vec{a}\cdot\vec{x})^2 - r^4)r_{,i} - 2r(\vec{a}\cdot\vec{x})a_i}{r^4 + (\vec{a}\cdot\vec{x})^2} \\ \ell_{i,j} &= \frac{(r^2\delta_{ij} - r\hat{a}_{ij} + a_i a_j)}{(r^3 + a^2 r)} + \frac{\left[(a^2r^2 - r^4)x_i + 2r^3(\vec{a}\times\vec{x})_i - (3r^2 + a^2)(\vec{a}\cdot\vec{x})a_i\right]r_{,j}}{(r^3 + a^2 r)^2}, \end{aligned}$$

where $\hat{a}_{ij} \equiv e_{ipj}a_p$.

A.3 Boosted Black Holes

A simple way to apply a boost transformation to a four-metric is to express the one-forms in the element in terms of the new coordinates. If we can calculate the boost in a preferred direction, this gets easier still. For the case of an originally spherically symmetric space-time, such as IEF, we choose the boost velocity v to be along the z-direction, and work in cylindrical polar coordinates: in the rest frame of the black hole, the coordinates are $X^\mu \equiv (T, Q, \Phi, Z)$[1], while in the lab frame (that is, the computational frame used for simulations), they are lower-case: $x^\mu = (t, q, \phi, z)$. In the rest frame of the black hole, the center is at $(0, 0, 0)$, while in the lab frame, it's at $(0, 0, vt)$. Note that $r \equiv \sqrt{q^2 + z^2}$, and $R \equiv \sqrt{Q^2 + Z^2}$. Under a boost in the z-direction, the rest-frame coordinates of the hole are related to the lab frame ones by:

$$\begin{aligned} T &= \gamma(t - vz) \ , \ Q = q \ , \ \Phi = \phi \ , \ Z = \gamma(z - vt) \\ \Rightarrow dT &= \gamma dt - v\gamma dz \ , \ dQ = dq \ , \ d\Phi = d\phi \ , \ dZ = \gamma dz - v\gamma dt \end{aligned}$$

In the boosted frame, the IEF metric is

$$\begin{aligned} ds^2 &= -dT^2 + dQ^2 + Q^2 d\Phi^2 + dZ^2 + \frac{2M}{R}\left(dT + \frac{Q}{R}dQ + \frac{Z}{R}dZ\right)^2 \\ &= -\gamma^2(dt - vdz)^2 + dq^2 + q^2 d\phi^2 + \gamma^2(dz - vdt)^2 \\ &\quad + \frac{2M}{R}\left[\gamma(dt - vdz) + \frac{q}{R}dq + \gamma\frac{Z}{R}(dz - vdt)\right]^2 \\ &= -dt^2 + dq^2 + q^2 d\phi^2 + dz^2 \\ &\quad + \frac{2M}{R}\left[\gamma\left(1 - \frac{vZ}{R}\right)dt + \frac{q}{R}dq + \gamma\left(\frac{Z}{R} - v\right)dz\right]^2 \\ &= -dt^2 + dq^2 + q^2 d\phi^2 + dz^2 + \frac{2M}{R}\left[\frac{\rho}{R}dt + \frac{q}{R}dq + \frac{\zeta}{R}dz\right]^2 . \end{aligned}$$

This implies a four-metric given by $g_{\mu\nu} = \eta_{\mu\nu} + 2H\ell_\mu\ell_\nu$, where

$$H = \frac{M}{R} \ , \ \ell_\mu = (\ell_t, \ell_q, \ell_\phi, \ell_z) = \left(\frac{\rho}{R}, \frac{q}{R}, 0, \frac{\zeta}{R}\right),$$

[1] Q here should not be confused with the densitized lapse discussed in the main text.

and $\rho \equiv \gamma(R - vZ)$, $\zeta \equiv \gamma(Z - vR)$. Note that

$$\begin{aligned} q^2 + \zeta^2 &= R^2 - Z^2 + \gamma^2(Z^2 - 2vZR + v^2R^2) \\ &= \gamma^2\left((1-v^2)(R^2 - Z^2) + Z^2 - 2vZR + v^2R^2\right) \\ &= \gamma^2(R^2 - 2vZR + v^2Z^2) \\ &= \rho^2. \end{aligned}$$

As ℓ_μ is still null, the inverse four-metric is very simply $g^{\mu\nu} = \eta^{\mu\nu} - 2H\ell^\mu\ell^\nu$. Now the three-metric is

$$(\gamma_{ij}) = \begin{pmatrix} 1 + \frac{2Mq^2}{R^3} & 0 & \frac{2Mq\zeta}{R^3} \\ 0 & q^2 & 0 \\ \frac{2Mq\zeta}{R^3} & 0 & 1 + \frac{2M\zeta^2}{R^3} \end{pmatrix},$$

the lapse function is

$$\alpha = \frac{1}{\sqrt{1 + \frac{2M\rho^2}{R^3}}},$$

the shift vector is

$$\left(\beta^q, \beta^\phi, \beta^z\right) = \frac{2M\rho}{R^3 + 2M\rho^2}\left(q, 0, \zeta\right),$$

and the extrinsic curvature is:

$$\begin{aligned} K_{qq} &= -\frac{M\left[2Mq^2\rho + \gamma R^2q^2 + \rho R(q^2 - 2Z^2)\right]}{R^4\sqrt{R^4 + 2MR\rho^2}} \\ K_{qz} &= -\frac{Mq\left[2M\rho\zeta + 3\gamma^2Zq^2 + \gamma(3Z^2 + r^2)\zeta\right]}{R^4\sqrt{R^4 + 2MR\rho^2}} \\ K_{\phi\phi} &= \frac{2Mq^2\rho}{R\sqrt{R^4 + 2MR\rho^2}} \\ K_{zz} &= \frac{M\left[(R-2M)\rho\zeta^2 - R\left(R\gamma\zeta^2 + 2\gamma R\rho^2 - 4q^2\gamma^2\rho\right)\right]}{R^4\sqrt{R^4 + 2MR\rho^2}} \end{aligned}$$

Then the three-metric determinant is

$$|\gamma| = 1 + \frac{2M\rho^2}{R^3},$$

and the trace of the extrinsic curvature is

$$K = \frac{M\rho^2\left[6M\rho + \gamma R(2R + 3vZ)\right]}{R(R^3 + 2M\rho^2)\sqrt{R(R^3 + 2M\rho^2)}}$$

To linear order in v, this can be approximated as:

$$K \approx \frac{2M(r+3M)}{\sqrt{r^3(r+2M)^3}} - v\frac{Mz(r+8M)}{\sqrt{r^3(r+2M)^5}}$$

Having calculated these, we can transform back to Cartesian coordinates. The Cartesian shift vector is

$$(\beta^x, \beta^y, \beta^z) = \frac{2M\rho}{R^3 + 2M\rho^2}\,(x, y, \zeta)\,,$$

In practice, the boost transformation will be calculated numerically using the boost matrix. However, the above derivation has proven useful in certain semi-analytic calculations.

A.4 Conformal Binary Black Hole Data

The only known exact solutions for a binary black hole system are for very special systems on a single time slice. These are, in general, conformally flat data, natural extensions of the Isotropic Schwarzschild solution (Section A.1.2).

A.4.1 Misner and Brill-Lindquist Data

The main conformal data schemes for two (or more) black holes are due to Misner [94] and Brill & Lindquist [33]. In both cases, the three-metric and extrinsic curvature are of the form (in Cartesian coordinates):

$$\begin{aligned} \gamma_{ij} &= \psi^4 \delta_{ij}, \\ K_{ij} &= 0. \end{aligned}$$

As the extrinsic curvature is zero, the momentum constraint is trivially satisfied with this scheme, while the Hamiltonian constraint reduces to the Laplace equation in flat space. For data representing initially very widely separated holes with masses m_A and m_B, this is solved by a conformal factor ψ given by

$$\psi = \left(1 + \frac{m_A}{2r_A} + \frac{m_B}{2r_B}\right)^4.$$

For finite separations, the topology of the data scheme becomes relevant, and this is where Misner and Brill-Lindquist data differ.

With Brill-Lindquist data, N holes are represented by $N+1$ sheets connected by N throats. The single common sheet represents our external universe, while the others represent the interior of each hole. Then the Hamiltonian constraint can be satisfied using a conformal factor[2]:

$$\psi = 1 + \sum_i^N \frac{\alpha_i}{r_i}.$$

Again, in the limit of infinite separation, the coefficients α_i become the bare masses of the N individual holes.

In contrast, Misner chose a two-sheeted topology – N holes are represented by N throats connecting an upper sheet (our external universe) to a lower sheet. The requirement that the data be reflection-symmetric across *every* hole means that the Brill-Lindquist form above will not work. Misner used a method-of-images approach to form a convergent infinite sum for ψ. This ψ still satisfied the Hamiltonian constraint.

[2] In fact, Brill & Lindquist allowed for charged holes also, necessitating *two* conformal factors

For both topologies[3], the solutions are time-symmetric, and apply only on the initial time slice. To find the complete time history of the holes requires numerical evolution.

A.4.2 Bowen-York data

One of the consequences of the time-symmetry of Misner and Brill-Lindquist data is that the holes cannot be moving. With this limitation, one could only attempt the head-on collision of holes with finite separation and no initial momentum.

In 1980, Bowen & York [28] found a general solution of the momentum constraint for conformally flat data, allowing the inclusion of both linear and angular momentum (P_i and S_i, respectively). For a single hole centered at the origin, the Bowen-York three-metric and extrinsic curvature are:

$$\begin{aligned}
\gamma_{ij} &= \psi^4 \delta_{ij}, \\
K_{ij} &= \psi^{-2} \left(\hat{K}_{Pij} + \hat{K}_{Sij} \right), \\
\hat{K}_{Pij} &= \frac{3}{2r^2} \left[P_i n_j + P_j n_i - (\delta_{ij} - n_i n_j) P^k n_k \right] \\
&\quad \mp \frac{3a^2}{2r^4} \left[P_i n_j + P_j n_i + (\delta_{ij} - 5 n_i n_j) P^k n_k \right], \\
\hat{K}_{Sij} &= \frac{3}{r^3} \left[\epsilon_{kim} J^m n^k n_j + \epsilon_{kjm} J^m n^k n_i \right].
\end{aligned}$$

(The second term on the right of $\hat{K}_{Pij}$ is only necessary for cases where one wishes to preserve the inversion symmetry found in Isotropic Schwarzschild, as with Misner data). As the three-metric is conformally flat, the Hamiltonian constraint takes a Poisson-like form:

$$\Delta \psi + \frac{1}{8} \psi^{-7} \hat{K}_{ij} \hat{K}^{ij} = 0$$

The nice thing about Bowen-York data is that, since the momentum constraint is linear, we can supply N holes with linear and angular momentum simply by forming a linear sum of such K_{ij}, one for each hole. In particular, boosts and spins can be added to Brill-Lindquist and Misner data sets[4]. As one only ever has to solve one elliptic equation – the Hamiltonian constraint – the initial data problem is very numerically tractable.

However, it should be noted that even if there is only one hole present, the Hamiltonian constraint will not be satisfied unless the extrinsic curvature vanishes. That is, the single-hole limit of this data is not a clean boosted and / or spinning black hole: there will be constraint violations and additional gravitational radiation present. In general, to produce

[3] Misner[93] gave another possible topology, containing one sheet and a wormhole. However, this is more restrictive, allowing exactly two holes, of equal mass.

[4] The Bowen-York extrinsic curvature for Misner-like data require some additional terms to satisfy the inversion symmetry at the throats.

a steady-state hole, we must solve the Hamiltonian constraint and evolve the resulting data until the excess radiation has left the system.

A.5 Linearized Wave Solutions

A.5.1 Linear Plane Waves

Moving away from black-hole solutions, a simple and useful approximate solution to the Einstein equations is that of a linear plane wave. Assuming a wave propagating along the z-axis on a flat background, the four-metric will be (in Cartesian coordinates):

$$ds^2 = -dt^2 + (1 - \epsilon U(z - vt))\, dx^2 - 2\epsilon V(z - vt) dx\, dy + (1 + \epsilon U(z - vt))\, dy^2 + dz^2.$$

The two fields U and V are independent wave functions – they could be thought of as the two polarizations of an arbitrary gravitational-wave signal. As this is a perturbation of flat space, we neglect terms of greater than first order in ϵ. Then the natural 3+1 decomposition is:

$$\begin{aligned} \alpha &= 1, \\ \beta^i &= 0, \\ \gamma_{ij} &= \begin{pmatrix} 1 - \epsilon U & -\epsilon V & 0 \\ -\epsilon V & 1 + \epsilon U & 0 \\ 0 & 0 & 1 \end{pmatrix}, \\ K_{ij} &= -\frac{1}{2} \begin{pmatrix} \epsilon v U' & \epsilon v V' & 0 \\ \epsilon v V' & -\epsilon v U' & 0 \\ 0 & 0 & 0 \end{pmatrix}. \end{aligned} \tag{A.8}$$

A.5.2 Spherical Plane Waves

A similar linear solution can be envisaged for a wave form moving radially inward or outward from the center of coordinates of the system. The corresponding four-metric would be (in spherical polar coordinates):

$$\begin{aligned} ds^2 &= -dt^2 + dr^2 + \left(1 - \frac{\epsilon U(r - vt)}{r}\right) r^2\, d\theta^2 - 2\frac{\epsilon V(r - vt)}{r} r^2\, \sin\theta\, d\theta\, d\phi \\ &\quad + \left(1 + \frac{\epsilon U(r - vt)}{r}\right) r^2\, \sin^2\theta\, d\phi^2 \end{aligned}$$

Again, this has a natural 3+1 decomposition; staying in spherical coordinates, this will be:

$$\begin{aligned}
\alpha &= 1, \\
\beta^i &= 0, \\
\gamma_{ij} &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & r^2 - \epsilon r U(r-vt) & -\epsilon r V(r-vt)\sin\theta \\ 0 & -\epsilon r V(r-vt)\sin\theta & r^2\sin^2\theta + \epsilon r U(r-vt)\sin^2\theta \end{pmatrix}, \\
K_{ij} &= -\frac{1}{2}\begin{pmatrix} 0 & 0 & 0 \\ 0 & \epsilon v r U' & \epsilon v r V'\sin\theta \\ 0 & \epsilon v r V'\sin\theta & -\epsilon v r U'\sin^2\theta \end{pmatrix}. \qquad \text{(A.9)}
\end{aligned}$$

Appendix B

Diagnostic Tools – Selected Details

This appendix draws together some extra details about the calculation and testing of certain diagnostic quantities introduced in Chap 3.

B.1 ADM Energy and Momentum

Below I present some numerical results on the accuracy of the ADM calculations for the test case of stationary black holes.

B.1.1 IEF and Kerr-Schild

First the formula was tested on IEF (Ingoing Eddington-Finkelstein). As it happens, for this space-time, the analytic integral converges exactly at *any* radius.

Results are presented at different resolutions and distances. The extreme distances were achieved by reducing the mass parameter M for a certain domain size, so that for instance, $M = 0.25$ and outer boundary at ± 30 really means outer boundary at $\pm 120M$. Results presented here have been renormalized to $M = 1$.

We note that the surface distance is irrelevant here – the error is the same for the same resolution. This is no doubt due to the behavior of the integrand for IEF – the radial dependence cancels exactly. The angular resolution does have an effect, though – higher

Δ	n_θ,n_ϕ	*x_outer*	*% error*
$0.8M$	49,98	$80M$	2.76
$0.4M$	49,98	$40M$	2.76
$0.4M$	97,194	$40M$	1.36
$0.4M$	25,50	$40M$	5.71
$0.2M$	49,98	$20M$	2.76
$0.2M$	25,50	$20M$	5.71
$0.1M$	49,98	$10M$	2.76
$0.05M$	49,98	$5M$	2.76

Table B.1: ADM energy for IEF solution in octant symmetry.

Δ	*x_outer*	*% error*
$0.8M$	$80M$	2.76
$0.4M$	$40M$	2.76
$0.2M$	$20M$	2.77
$0.1M$	$10M$	2.77
$0.05M$	$5M$	2.78

Table B.2: Same as Table 1, but for Kerr-Schild with $a = 0.5M$.

resolution means lower error, as expected. We also see that the method of using smaller mass parameters instead of more grid points works fine. This is observed in later results, so repetitions will be omitted.

The errors again seem to be identical for all distances, as for IEF. Accordingly, later results will be presented with an "adjusted" column, featuring the percentage error less that error attributable to finite-differencing (2.76% for medium angular resolution), in the hope of leaving only a distance-dependent error measure.

B.1.2 Isotropic Schwarzschild

The same test was carried out for Isotropic Schwarzschild. This does *not* have the same nice properties as IEF: the finite-R integral does not reduce exactly to M:

$$\begin{aligned}
\gamma_{ij} &= \left(1+\frac{M}{2r}\right)^4 \delta_{ij} \\
\Rightarrow \gamma_{ij,k} &= \left(1+\frac{M}{2r}\right)^3 \delta_{ij}\left(-\frac{2Mx_k}{r^3}\right) \\
\Rightarrow \gamma_{ij,i} &= \left(1+\frac{M}{2r}\right)^3 \left(-\frac{2Mx_j}{r^3}\right), \quad \gamma_{ii,j} = \left(1+\frac{M}{2r}\right)^3 \left(-\frac{6Mx_j}{r^3}\right) \\
\Rightarrow I_{trad} &= \frac{1}{16\pi}\oint_{r=R}\left(1+\frac{M}{2r}\right)^3 \left(\frac{4M}{r^2}\right) r^2 d\Omega \\
&= \frac{1}{16\pi} 4M \left(1+\frac{M}{2R}\right)^3 4\pi = M\left(1+\frac{M}{2R}\right)^3 \approx M\left(1+\frac{3M}{2R}\right)
\end{aligned}$$

As the ADM integrand has a non-trivial dependence on the radial position, it's not surprising that the accuracy declines dramatically as one calculates the integral at smaller and smaller R. The "adjusted" error seems to decline as R^{-1}, at least for $R > 10M$ or so.

B.1.3 ADM Momentum

The linear momentum in a spatial slice can be calculated in a manner similar to the energy:

Δ	*x_outer*	*% error*	*adjusted % error*	*expected % error*
$0.8M$	$80M$	4.70	1.94	1.87
$0.4M$	$40M$	6.67	3.91	3.80
$0.2M$	$20M$	10.67	7.91	7.69
$0.2M$	$10M$	18.97	16.21	15.76
$0.2M$	$5M$	36.81	34.05	33.10

Table B.3: Isotropic Schwarzschild.

x_outer	*% M error*	*% E error*	*Pz error*	*falloff error*	*adjusted error*
$40M$	4.366	4.362	5.650%	1.225%	4.425%
$20M$	4.366	4.362	6.462%	2.006%	4.456%
$10M$	4.366	4.362	7.435%	2.945%	4.490%
$5M$	4.366	4.362	8.737%	4.201%	4.536%
$10M$	2.132	2.131	5.126%	2.945%	2.181%
$10M$	4.366	4.360	6.189%	1.757%	4.432%
$10M$	4.366	4.357	4.176%	-0.164%	4.340%
$10M$	4.366	4.355	0.673%	-3.506%	4.179%
$10M$	4.366	4.351	-6.470%	-10.302%	3.832%

Table B.4: ADM z-momentum for boosted IEF solution in quadrant symmetry, with $n_\theta = 49$, $n_\phi = 26$.

$$\begin{aligned} J^i &\equiv \frac{1}{8\pi} \oint_{r=R} (K_{pi} N^p - K N_i)\, dA \\ P^i_{ADM} &= \lim_{R\to\infty} J^i. \end{aligned}$$

For angular momentum, in the z-direction, $N^p = x\delta^p_2 - y\delta^p_1$. Similarly, in the x-direction, $N^p = y\delta^p_3 - z\delta^p_2$ and in the y-direction, $N^p = z\delta^p_1 - x\delta^p_3$.

B.2 Weyl Scalars

Below I provide some more details on the calculation of the Weyl scalars defined in eqn (3.31), using the ADM quantities typically available in numerical simulations. The method follows [76], but uses the more standard metric signature $(-,+,+,+)$.

B.2.1 Derivation from 3+1 ADM Quantities

The 3+1 split already implies a unit normal vector, $\hat{n}^a$. Following [120], we can decompose the Weyl tensor into its "electric" and "magnetic" parts:

$$C_{abcd} = 4\hat{n}_{[a}E_{b][c}\hat{n}_{d]} + 2\varepsilon_{ab}{}^{e}B_{e[c}\hat{n}_{d]} + 2\varepsilon_{cd}{}^{e}B_{e[a}\hat{n}_{b]} + \varepsilon_{ab}{}^{e}\varepsilon_{cd}{}^{f}E_{ef}. \tag{B.1}$$

In terms of ADM quantities, E_{ab} and B_{ab} are given by

$$\begin{aligned} E_{ab} &\equiv -C_{acbd}\hat{n}^c\hat{n}^d = -R_{ab} + K_a{}^cK_{bc} - KK_{ab}, \\ B_{ab} &\equiv -{}^*C_{acbd}\hat{n}^c\hat{n}^d = -\varepsilon_a{}^{cd}D_cK_{db}, \end{aligned}$$

where ${}^*C_{abcd} \equiv \frac{1}{2}{}^4\varepsilon_{ab}{}^{mn}C_{mncd}$ is the dual to the Weyl tensor. We take an arbitrary orthonormal triad $(\hat{u}^a, \hat{v}^a, \hat{w}^a)$, orthogonal to $\hat{n}^a$, such that $\varepsilon_{abc}\hat{u}^a\hat{v}^b\hat{w}^c = 1$. Then

$$\varepsilon_{ab}{}^{c}\hat{u}^a\hat{v}^b = \hat{w}^c, \tag{B.2}$$

and so on, by cyclic symmetry in the spatial triad. The four-metric can be written in terms of these spatial vectors, plus the time-like normal:

$$g_{ab} = -\hat{n}_a\hat{n}_b + \hat{u}_a\hat{u}_b + \hat{v}_a\hat{v}_b + \hat{w}_a\hat{w}_b. \tag{B.3}$$

We can use these spatial vectors, together with the time-like $\hat{n}^a$, to define a null tetrad:

$$\begin{aligned} l^a &= \frac{1}{\sqrt{2}}\left(\hat{n}^a + \hat{u}^a\right), \\ n^a &= \frac{1}{\sqrt{2}}\left(\hat{n}^a - \hat{u}^a\right), \\ m^a &= \frac{1}{\sqrt{2}}\left(\hat{v}^a - i\hat{w}^a\right), \end{aligned}$$

Now we use this to describe the Weyl scalars in terms of the electric and magnetic

tensors. The most involved calculation is the "monopole" term, Ψ_2, which we do first:

$$
\begin{aligned}
\Psi_2 &\equiv C_{abcd} l^a m^b \bar{m}^c n^d \\
&= \frac{1}{2} C_{abcd} (\hat{n}^a \hat{n}^d - \hat{n}^a \hat{u}^d + \hat{u}^a \hat{n}^d - \hat{u}^a \hat{u}^d) m^b \bar{m}^c \\
&= \frac{1}{2} \Big[E_{bc} - E_{bd} \hat{u}^d \hat{n}_c - \varepsilon_{cd}{}^e \hat{u}^d B_{eb} + \hat{n}_b E_{ac} \hat{u}^a - \varepsilon_{ab}{}^e \hat{u}^a B_{ec} \\
&\quad - \hat{n}_b E_{ad} \hat{u}^a \hat{u}^d \hat{n}_c + \varepsilon_{ab}{}^e \hat{u}^a B_{ed} \hat{u}^d \hat{n}_c - \varepsilon_{cd}{}^e \hat{u}^d B_{ea} \hat{u}^a \hat{n}_b - \varepsilon_{ab}{}^e \hat{u}^a \varepsilon_{cd}{}^f \hat{u}^d E_{ef} \Big] m^b \bar{m}^c \\
&= \frac{1}{4} \Big[E_{bc} - \varepsilon_{cd}{}^e \hat{u}^d B_{eb} - \varepsilon_{ab}{}^e \hat{u}^a B_{ec} - \varepsilon_{ab}{}^e \hat{u}^a \varepsilon_{cd}{}^f \hat{u}^d E_{ef} \Big] \Big[(\hat{v}^b \hat{v}^c + \hat{w}^b \hat{w}^c) + i (\hat{v}^b \hat{w}^c - \hat{w}^b \hat{v}^c) \Big] \\
&= \frac{1}{4} \Big(E_{bc} \Big[(\hat{v}^b \hat{v}^c + \hat{w}^b \hat{w}^c) + i (\hat{v}^b \hat{w}^c - \hat{w}^b \hat{v}^c) \Big] \\
&\quad - B_{eb} \Big[(-\hat{w}^e + i \hat{v}^e) \hat{v}^b + (\hat{v}^e + i \hat{w}^e) \hat{w}^b \Big] \\
&\quad - B_{ec} \left[(\hat{w}^e + i \hat{v}^e) \hat{v}^c + (-\hat{v}^e + i \hat{w}^e) \hat{w}^c \right] - E_{ef} \Big[(-\hat{w}^f + i \hat{v}^f) \hat{w}^e - (\hat{v}^f + i \hat{w}^f) \hat{v}^e \Big] \Big) \\
&= \frac{1}{2} (E_{bc} - i B_{bc}) \Big[\hat{v}^b \hat{v}^c + \hat{w}^b \hat{w}^c \Big] \\
&= \frac{1}{2} (E_{bc} - i B_{bc}) \Big[g^{bc} + \hat{n}^b \hat{n}^c - \hat{u}^b \hat{u}^c \Big] \\
&= -\frac{1}{2} (E_{bc} - i B_{bc}) \hat{u}^b \hat{u}^c .
\end{aligned}
$$

In going from line 3 to line 4 above, we used the orthogonality of $\hat{n}^a$ and m^a; from line 4 to line 5, we used the cyclic relation (B.2); from line 5 to line 6, we used the symmetry of E_{ab} and B_{ab}; from line 6 to line 7, we used the decomposition of the four-metric (B.3); and from line 7 to line 8, we used the spatial and trace-free nature of E_{ab} and B_{ab}.

In a similar manner, we can treat the other scalars. The "transverse ingoing" scalar,

Ψ_0, becomes:

$$\begin{aligned}
\Psi_0 &\equiv C_{abcd}l^a m^b l^c m^d = -C_{abcd}l^a m^b m^c l^d \\
&= -\frac{1}{2}C_{abcd}(\hat{n}^a\hat{n}^d + \hat{n}^a\hat{u}^d + \hat{u}^a\hat{n}^d + \hat{u}^a\hat{u}^d)m^b m^c \\
&= -\frac{1}{2}\Big[E_{bc} + E_{bd}\hat{u}^d\hat{n}_c + \varepsilon_{cd}{}^e\hat{u}^d B_{eb} + \hat{n}_b E_{ac}\hat{u}^a - \varepsilon_{ab}{}^e\hat{u}^a B_{ec} + \hat{n}_b E_{ad}\hat{u}^a\hat{u}^d\hat{n}_c \\
&\quad -\varepsilon_{ab}{}^e\hat{u}^a B_{ed}\hat{u}^d\hat{n}_c + \varepsilon_{cd}{}^e\hat{u}^d B_{ea}\hat{u}^a\hat{n}_b + \varepsilon_{ab}{}^e\hat{u}^a\varepsilon_{cd}{}^f\hat{u}^d E_{ef}\Big]m^b m^c \\
&= -\frac{1}{2}\Big[E_{bc} + \varepsilon_{cd}{}^e\hat{u}^d B_{eb} - \varepsilon_{ab}{}^e\hat{u}^a B_{ec} + \varepsilon_{ab}{}^e\hat{u}^a\varepsilon_{cd}{}^f\hat{u}^d E_{ef}\Big]m^b m^c \\
&= -\frac{1}{2}\Big[E_{bc} + 2\varepsilon_{cd}{}^e\hat{u}^d B_{eb}\Big]m^b m^c - \frac{1}{4}\varepsilon_{ab}{}^e\hat{u}^a\varepsilon_{cd}{}^f\hat{u}^d E_{ef}(\hat{v}^b\hat{v}^c - i\hat{v}^b\hat{w}^c - i\hat{w}^b\hat{v}c - \hat{w}^b\hat{w}^c) \\
&= -\frac{1}{2}\Big[E_{bc} + 2\varepsilon_{cd}{}^e\hat{u}^d B_{eb}\Big]m^b m^c - \frac{1}{4}E_{ef}(-\hat{w}^e\hat{w}^f - i\hat{w}^e\hat{v}^f - i\hat{v}^e\hat{w}^f + \hat{v}^e\hat{v}^f) \\
&= -\frac{1}{2}\Big[E_{bc} + 2\varepsilon_{cd}{}^e\hat{u}^d B_{eb}\Big]m^b m^c - \frac{1}{2}E_{ef}m^e m^f \\
&= -\Big[E_{bc} + \varepsilon_{cd}{}^e\hat{u}^d B_{eb}\Big]m^b m^c
\end{aligned}$$

The "ingoing longitudinal" scalar, Ψ_1, becomes:

$$\begin{aligned}
\Psi_1 &\equiv C_{abcd}l^a n^b l^c m^d = C_{abcd}l^a m^b l^c n^d \\
&= \frac{1}{2}C_{abcd}(\hat{n}^a\hat{n}^d - \hat{n}^a\hat{u}^d + \hat{u}^a\hat{n}^d - \hat{u}^a\hat{u}^d)m^b l^c \\
&= \frac{1}{2}\Big[E_{bc} - E_{bd}\hat{u}^d\hat{n}_c - \varepsilon_{cd}{}^e\hat{u}^d B_{eb} + \hat{n}_b E_{ac}\hat{u}^a - \varepsilon_{ab}{}^e\hat{u}^a B_{ec} \\
&\quad -\hat{n}_b E_{ad}\hat{u}^a\hat{u}^d\hat{n}_c + \varepsilon_{ab}{}^e\hat{u}^a B_{ed}\hat{u}^d\hat{n}_c - \varepsilon_{cd}{}^e\hat{u}^d B_{ea}\hat{u}^a\hat{n}_b - \varepsilon_{ab}{}^e\hat{u}^a\varepsilon_{cd}{}^f\hat{u}^d E_{ef}\Big]m^b l^c \\
&= \frac{1}{2\sqrt{2}}\Big[E_{bc} - E_{bd}\hat{u}^d\hat{n}_c - \varepsilon_{cd}{}^e\hat{u}^d B_{eb} - \varepsilon_{ab}{}^e\hat{u}^a B_{ec} \\
&\quad +\varepsilon_{ab}{}^e\hat{u}^a B_{ed}\hat{u}^d\hat{n}_c - \varepsilon_{ab}{}^e\hat{u}^a\varepsilon_{cd}{}^f\hat{u}^d E_{ef}\Big]m^b(\hat{n}^c + \hat{u}^c) \\
&= \frac{1}{2\sqrt{2}}\Big[E_{bc}\hat{u}^c + E_{bd}\hat{u}^d - \varepsilon_{ab}{}^e\hat{u}^a B_{ec}\hat{u}^c - \varepsilon_{ab}{}^e\hat{u}^a B_{ed}\hat{u}^d\Big]m^b \\
&= \frac{1}{\sqrt{2}}\left[E_{bc}\hat{u}^c - \varepsilon_{ab}{}^e\hat{u}^a B_{ec}\hat{u}^c\right]m^b
\end{aligned}$$

The "outgoing longitudinal" scalar, Ψ_3, becomes:

$$\begin{aligned}
\Psi_3 &\equiv C_{abcd} l^a n^b \bar{m}^c n^d \\
&= \frac{1}{2} C_{abcd} (\hat{n}^a \hat{n}^d - \hat{n}^a \hat{u}^d + \hat{u}^a \hat{n}^d - \hat{u}^a \hat{u}^d) n^b \bar{m}^c \\
&= \frac{1}{2} \Big[E_{bc} - E_{bd} \hat{u}^d \hat{n}_c - \varepsilon_{cd}{}^e \hat{u}^d B_{eb} + \hat{n}_b E_{ac} \hat{u}^a - \varepsilon_{ab}{}^e \hat{u}^a B_{ec} \\
&\quad - \hat{n}_b E_{ad} \hat{u}^a \hat{u}^d \hat{n}_c + \varepsilon_{ab}{}^e \hat{u}^a B_{ed} \hat{u}^d \hat{n}_c - \varepsilon_{cd}{}^e \hat{u}^d B_{ea} \hat{u}^a \hat{n}_b - \varepsilon_{ab}{}^e \hat{u}^a \varepsilon_{cd}{}^f \hat{u}^d E_{ef} \Big] n^b \bar{m}^c \\
&= \frac{1}{2\sqrt{2}} \Big[E_{bc} - \varepsilon_{cd}{}^e \hat{u}^d B_{eb} + \hat{n}_b E_{ac} \hat{u}^a - \varepsilon_{ab}{}^e \hat{u}^a B_{ec} \\
&\quad - \varepsilon_{cd}{}^e \hat{u}^d B_{ea} \hat{u}^a \hat{n}_b - \varepsilon_{ab}{}^e \hat{u}^a \varepsilon_{cd}{}^f \hat{u}^d E_{ef} \Big] (\hat{n}^b - \hat{u}^b) \bar{m}^c \\
&= \frac{1}{\sqrt{2}} \Big[-E_{bc} \hat{u}^b + \varepsilon_{cd}{}^e \hat{u}^d B_{eb} \hat{u}^b \Big] \bar{m}^c
\end{aligned}$$

Finally, the "outgoing transverse" scalar, Ψ_4, becomes:

$$\begin{aligned}
\Psi_4 &\equiv C_{abcd} n^a \bar{m}^b n^c \bar{m}^d = -C_{abcd} n^a \bar{m}^b \bar{m}^c n^d \\
&\quad - \frac{1}{2} C_{abcd} (\hat{n}^a \hat{n}^d - \hat{n}^a \hat{u}^d - \hat{u}^a \hat{n}^d + \hat{u}^a \hat{u}^d) \bar{m}^b \bar{m}^c \\
&= -\frac{1}{2} \Big[E_{bc} - E_{bd} \hat{u}^d \hat{n}_c - \varepsilon_{cd}{}^e \hat{u}^d B_{eb} - \hat{n}_b E_{ac} \hat{u}^a + \varepsilon_{ab}{}^e \hat{u}^a B_{ec} + \hat{n}_b E_{ad} \hat{u}^a \hat{u}^d \hat{n}_c \\
&\quad - \varepsilon_{ab}{}^e \hat{u}^a B_{ed} \hat{u}^d \hat{n}_c + \varepsilon_{cd}{}^e \hat{u}^d B_{ea} \hat{u}^a \hat{n}_b + \varepsilon_{ab}{}^e \hat{u}^a \varepsilon_{cd}{}^f \hat{u}^d E_{ef} \Big] \bar{m}^b \bar{m}^c \\
&= -\frac{1}{2} \Big[E_{bc} - \varepsilon_{cd}{}^e \hat{u}^d B_{eb} + \varepsilon_{ab}{}^e \hat{u}^a B_{ec} + \varepsilon_{ab}{}^e \hat{u}^a \varepsilon_{cd}{}^f \hat{u}^d E_{ef} \Big] \bar{m}^b \bar{m}^c \\
&= -\frac{1}{2} \Big[E_{bc} - 2\varepsilon_{cd}{}^e \hat{u}^d B_{eb} \Big] \bar{m}^b \bar{m}^c - \frac{1}{4} E_{ef} (-\hat{w}^e \hat{w}^f + i\hat{w}^e \hat{v}^f + i\hat{v}^e \hat{w}^f + \hat{v}^e \hat{v}^f) \\
&= -\frac{1}{2} \Big[E_{bc} - 2\varepsilon_{cd}{}^e \hat{u}^d B_{eb} \Big] \bar{m}^b \bar{m}^c - \frac{1}{2} E_{ef} \bar{m}^e \bar{m}^f \\
&= -\Big[E_{bc} - \varepsilon_{cd}{}^e \hat{u}^d B_{eb} \Big] \bar{m}^b \bar{m}^c
\end{aligned}$$

These results can be tidied up somewhat by projecting the electric and magnetic spatial tensors along and perpendicular to the preferred spatial direction, $\hat{u}^a$. Define

$$\begin{aligned}
e &\equiv E_{ab} \hat{u}^a \hat{u}^b, \\
e_a &\equiv E_{bc} \hat{u}^b \left(\delta^c_a - \hat{u}_a \hat{u}^c \right) \\
&= E_{ab} \hat{u}^b - e \hat{u}_a, \\
e_{ab} &\equiv E_{cd} \left(\delta^c_a - \hat{u}_a \hat{u}^c \right) \left(\delta^d_b - \hat{u}_b \hat{u}^d \right) + \frac{1}{2} e \left(\gamma_{ab} - \hat{u}_a \hat{u}_b \right),
\end{aligned}$$

In terms of these quantities, E_{ab} and B_{ab} can be reconstructed as:

$$\begin{aligned} E_{ab} &= \frac{1}{2}\left(3\hat{u}_a\hat{u}_b - \gamma_{ab}\right)e + 2e_{(a}\hat{u}_{b)} + e_{ab} \\ B_{ab} &= \frac{1}{2}\left(3\hat{u}_a\hat{u}_b - \gamma_{ab}\right)b + 2b_{(a}\hat{u}_{b)} + b_{ab}. \end{aligned}$$

Defining a "rotation operator" $J_a^b \equiv \varepsilon_a{}^{bcd}\hat{u}_c\hat{n}_d = \varepsilon_a{}^{bc}\hat{u}_c$, the Weyl scalars can also be written as:

$$\begin{aligned} \Psi_0 &= -\left(e_{ab} - J_a^c b_{bc}\right)m^a m^b \\ \Psi_1 &= \frac{1}{\sqrt{2}}\left(e_a + J_a^c b_c\right)m^a \\ \Psi_2 &= -\frac{1}{2}\left(e - ib\right) \\ \Psi_3 &= -\frac{1}{\sqrt{2}}\left(e_a + J_a^c b_c\right)\bar{m}^a \\ \Psi_4 &= -\left(e_{ab} + J_a^c b_{bc}\right)\bar{m}^a\bar{m}^b \end{aligned}$$

B.2.2 Choice in Triads and Tetrads

The above requires that we have a tetrad available to us. Apart from $\hat{n}^a$, what directions to choose is somewhat arbitrary. Many perturbative studies of back hole space-times are done using the Kinnersley tetrad. In general, however, our space-time will be far from stationary, and may not have a well-defined mass or spin, both of which are required for the Kinnersley tetrad. In general, using the simpler "radial" tetrad (basically choosing the radial and angular directions) should be "as good". Whatever triad is chosen, we orthonormalize them via a Gram-Schmidt process. This procedure destroys the democracy of the group, however, so it is good to know which unit vector we'd like to preserve the most. From the formalism above, this is what we call $\hat{u}^a$. For the radial and Kinnersley tetrad cases, this should be the radial unit vector.

The simple "radial" tetrad is defined by choosing the unit radial vector as the preferred spatial vector $\hat{u}^a$, with the orthogonal (in flat space) angular directions for the other two spatial vectors:

$$\begin{aligned} \hat{u}^a &= [0,1,0,0] &\leftrightarrow& \quad \tfrac{1}{r}[0,x,y,z], \\ \hat{v}^a &= r^{-1}[0,0,1,0] &\leftrightarrow& \quad (qr)^{-1}\left[0,xz,yz,-q^2\right], \\ \hat{w}^a &= q^{-1}[0,0,0,1] &\leftrightarrow& \quad q^{-1}[0,-y,x,0], \end{aligned}$$

where the first arrays are the spherical polar coordinate expressions, and the second the Cartesian expressions.

Alternatively, if we view deviations from Kerr to be perturbations, we can use something like the Kinnersley tetrad (*à la* Teukolsky). In Boyer-Lindquist coordinates (which reduce to the Schwarzschild form for zero spin — see section A.2.1), the Kinnersley tetrad is (spherical components):

$$\begin{aligned}
\ell^a &= \left[\frac{r^2+a^2}{\Delta}, 1, 0, \frac{a}{\Delta}\right] \\
n^a &= \left[\frac{r^2+a^2}{2\Sigma}, -\frac{\Delta}{2\Sigma}, 0, \frac{a}{2\Sigma}\right] \\
m^a &= \left[\frac{ia\sin\theta}{\sqrt{2}(r+ia\cos\theta)}, 0, \frac{1}{\sqrt{2}(r+ia\cos\theta)}, \frac{i}{\sqrt{2}(r+ia\cos\theta)\sin\theta}\right].
\end{aligned}$$

Campanelli et al [40] adapt this for horizon-penetrating (Kerr-Schild) coordinates by a simple rescaling of ℓ^a and m^a to avoid singular behavior at the horizon, though their rescaled tetrad is dimensionally unbalanced, in that $\hat{n}^a$ can't be constructed by simply adding them together. If we correct the dimensions by multiplying and dividing by r^2, we get:

$$\begin{aligned}
\ell^a &= \left[\frac{\Delta+4Mr}{r^2}, \frac{\Delta}{r^2}, 0, \frac{2a}{r^2}\right] \\
n^a &= \left[\frac{r^2}{2\Sigma}, -\frac{r^2}{2\Sigma}, 0, 0\right] \\
m^a &= \left[\frac{ia\sin\theta}{\sqrt{2}(r+ia\cos\theta)}, 0, \frac{1}{\sqrt{2}(r+ia\cos\theta)}, \frac{i}{\sqrt{2}(r+ia\cos\theta)\sin\theta}\right].
\end{aligned}$$

However, this yields an orthonormal tetrad with mixed time- and space-components. We can obtain cleaner separation by assuming the spatial four-vectors have zero time component. Then some experimentation with Maple suggests the following:

$$\begin{aligned}
\hat{n}^a &= \tfrac{1}{\sqrt{\Sigma(\Sigma+2Mr)}}\left[\Sigma+2Mr, -2Mr, 0, 0\right] &\rightarrow& \left[\sqrt{\tfrac{r+2M}{r}}, -\tfrac{2M}{\sqrt{r^2+2Mr}}, 0, 0\right], \\
\hat{u}^a &= \sqrt{\tfrac{\Sigma}{\Sigma+2Mr}}\left[0, 1, 0, 0\right] &\rightarrow& \left[0, \sqrt{\tfrac{r}{r+2M}}, 0, 0\right], \\
\hat{v}^a &= \tfrac{1}{\sqrt{\Sigma}}\left[0, 0, 1, 0\right] &\rightarrow& \left[0, 0, \tfrac{1}{r}, 0\right], \\
\hat{w}^a &= \tfrac{1}{\sqrt{\Sigma}}\left[0, a\sin\theta, 0, \tfrac{1}{\sin\theta}\right] &\rightarrow& \left[0, 0, 0, \tfrac{1}{q}\right],
\end{aligned}$$

where the second set of vectors shown are in the Schwarzschild (zero-spin) limit. These have the advantage that, due to the zero time components, if the spatial four-vectors are orthonormal, so are the three-vectors corresponding to their spatial parts. Additionally, this reduces to the radial tetrad when mass and spin are set to zero.

B.2.3 Numerical Tests of the Weyl Scalars

It is difficult to come up with good test cases for the numerical calculation of the Weyl scalars. The standard black-hole solutions – with a sensible choice of tetrad (e.g., the radial or adapted Kinnersley above) – will possess only one non-trivial component, $Re(\Psi_2)$. To give more confidence in the "radiation" scalars, Ψ_0 and Ψ_4, I consider a linearized plane wave traveling along the z-axis on a flat (Minkowski) background (see [42], p. 521). Here, a Cartesian tetrad is better adapted to the extraction of radiation than the spherical-polar tetrads discussed above. Take a four-metric of

$$ds^2 = -dt^2 + (1 - \epsilon U(z - vt))\, dx^2 - 2\epsilon V(z - vt) dx\, dy + (1 + \epsilon U(z - vt))\, dy^2 + dz^2$$

Then, to linear order in ϵ, the only non-zero Weyl scalars will be:

$$\begin{aligned} \Psi_0 &= \frac{1}{4}\epsilon(1 - v)^2 \left[U'' + iV''\right], \\ \Psi_4 &= \frac{1}{4}\epsilon(1 + v)^2 \left[U'' - iV''\right], \end{aligned}$$

where the prime indicates differentiation with respect to the argument $(z - vt)$. In particular, for left-traveling gravitational waves, the wave velocity is $v = -1$ (that is, the wave argument is $z + t$). Then by direct calculation, the only non-zero Weyl scalar is

$$\Psi_0 = \epsilon U''(z + t) + i\epsilon V''(z + t),$$

while for right-traveling waves, the wave velocity is $v = 1$, and the only non-zero Weyl scalar is

$$\Psi_4 = \epsilon U''(z - t) - i\epsilon V''(z - t),$$

This linearized solution was used to test the correctness of the Ψ_0 and Ψ_4 calculation procedures of the code for small amplitudes ($\epsilon \sim 10^{-5}$ and smaller).

B.3 Quasi-normal Mode Fourier Analysis

When looking at the gravitational-wave profile in the far-field region of a ringing-down black hole, a Fourier transform should contain several well-resolved peaks at the characteristic QNM frequencies for that hole.

Taking a time signal of the form $s(t) = A\exp\left(i2\pi ft - t/\tau\right)$ (for $t > 0$ only – the signal will not be eternal), the FT will be, up to arbitrary normalization:

$$\begin{aligned}
\tilde{s}(\nu) &= \int_{\infty}^{\infty} s(t)e^{-i2\pi\nu t}dt \\
&= A\int_{0}^{\infty} e^{i2\pi(f-\nu)t - t/\tau}dt \\
&= \frac{A}{-1/\tau + i2\pi(f-\nu)}\int_{0}^{\infty} e^{u}du \\
&= \frac{A\tau}{1 - i2\pi(f-\nu)\tau} = \frac{A[1+i2\pi(f-\nu)]\tau}{1+4\pi^2(f-\nu)^2\tau^2} \\
\Rightarrow Re\left[\tilde{s}(\nu)\right] &= \frac{A\tau}{1+4\pi^2(f-\nu)^2\tau^2}
\end{aligned}$$

The real part of the transform will have its maximum $\tilde{s}_{max} = A\tau$ at $\nu^* = f$, as we would expect. The α-peak, where $\tilde{s}$ is a fraction α of the maximum, will be achieved at frequencies ν_α given by

$$\begin{aligned}
\frac{A\tau}{1+4\pi^2(f-\nu_\alpha)^2\tau^2} &= \alpha A\tau \\
\Rightarrow 1+4\pi^2(f-\nu_\alpha)^2\tau^2 &= \frac{1}{\alpha} \\
\Rightarrow \nu_\alpha - f &= \pm\frac{1}{2\pi\tau}\sqrt{\frac{1}{\alpha}-1} \\
\Rightarrow \nu_\alpha &= f \pm \frac{1}{2\pi\tau}\sqrt{\frac{1}{\alpha}-1}.
\end{aligned}$$

For instance, the half-width occurs at $\nu_{\frac{1}{2}} = f \pm \frac{1}{2\pi\tau}$.

Sources of error include the effect of neighboring peaks on each other – both the apparent width and the peak center itself will be affected. To estimate the effect of one peak on another, consider a Fourier spectrum containing two normal frequencies only, at frequencies $\nu = (f_1, f_2)$, where $f_2 - f_1 = \Delta > 0$. Assume further, for ease of computation, that the two modes have the same amplitude A and damping time τ[1] Then the combined transform will

[1]This is quite reasonable for black-hole QNMs – the $n = 1$ modes' damping times quickly asymptote to a value of $\tau \approx 10M$, as can be seen in Table 3.1.

look like:

$$\begin{aligned}
\tilde{s}(\nu) &= A\tau\left[\frac{1}{1+4\pi^2(\nu-f_1)^2\tau^2}+\frac{1}{1+4\pi^2(f_2-\nu)^2\tau^2}\right] \\
\Rightarrow \frac{d\tilde{s}}{d\nu} &= -8A\pi^2\tau^3\left[\frac{(\nu-f_1)}{(1+4\pi^2(\nu-f_1)^2\tau^2)^2}-\frac{(f_2-\nu)}{(1+4\pi^2(f_2-\nu)^2\tau^2)^2}\right] \\
&= 0 \text{ at } \nu = \nu* \\
\Rightarrow \frac{(\nu^*-f_1)}{(1+4\pi^2(\nu^*-f_1)^2\tau^2)^2} &= \frac{(f_2-\nu^*)}{(1+4\pi^2(f_2-\nu^*)^2\tau^2)^2}
\end{aligned}$$

These frequencies ν^* define the observed peaks. We are looking for the effect of the presence of mode 2 on the FT position of mode 1's peak. Say that the observed peak lies at a value $\nu^* = f_1 + \epsilon$. Then

$$\begin{aligned}
\epsilon\left[1+4\pi^2(\Delta-\epsilon)^2\tau^2\right]^2 &= (\Delta-\epsilon)\left[1+4\pi^2\epsilon^2\tau^2\right]^2 \\
\Rightarrow \epsilon\left(1+4\pi^2\Delta^2\tau^2\right)^2\left[1-\frac{8\pi^2\tau^2\Delta\epsilon}{1+4\pi^2\Delta^2\tau^2}+O(\epsilon^2)\right] &= (\Delta-\epsilon)\left[1+O(\epsilon^2)\right] \\
\Rightarrow \epsilon\left(1+4\pi^2\Delta^2\tau^2\right)^2 &= (\Delta-\epsilon)1+O(\epsilon^2) \\
\Rightarrow \epsilon &\approx \frac{\Delta}{1+(1+4\pi^2\Delta^2\tau^2)^2}
\end{aligned}$$

To get an idea of the size of this correction to the peak position, we take a few standard values. For fundamental modes, the damping times are $\tau \approx 10M$ (see Table 3.1). Also, $Re(\omega)$ between even-ℓ modes is $\sim 0.4M^{-1}$. Then the (dimensionless) denominator here is ≈ 300, while the numerator is $\approx 0.064M^{-1}$. This implies a peak shift $\epsilon \approx 0.00025M^{-1}$, or around 0.2% of the $\ell = 4$ value. That is, for slowly-damped fundamental modes, this effect will be negligible. It is possible, however, that the peak shift could be considerable for other mode combinations, perhaps with non-fundamental QNMs or other modes altogether.

B.4 Proper Separation

The initial data used in the binary mergers is parametrized by the individual black hole masses, the boosts given to them, and their initial separation. This last, however, is only the Euclidean separation – the coordinate separation on our computational grid. A more meaningful quantity would be the proper separation of the holes. This can be calculated numerically during evolution, but for the simple initial data sets used here, it should be possible to estimate the proper separation analytically.

To calculate the proper distance between apparent horizons, we need to integrate along the path joining them. In our case, as the holes are separated along the z-axis only, we know that the three-dimensional line-element is $ds_3 = \sqrt{\gamma_{zz}}dz$. Assume that hole A is centered at $z = -d/2$, and hole B at $z = d/2$, and that the apparent horizons have radius ℓ. Then the horizon-to-horizon proper separation D of the holes is:

$$D = \int_{z=-d/2+\ell}^{d/2-\ell} \sqrt{\gamma_{zz}}\, dz = 2\int_{z=0}^{d/2-\ell} \sqrt{\gamma_{zz}}\, dz$$

where I have used the symmetry of the initial holes to simplify the integral. Now, for boosted binary HuMaSh data,

$$\gamma_{zz} = 1 + \frac{2\, m_A\, \zeta_A^2}{R_A^3} + \frac{2\, m_B\, \zeta_B^2}{R_B^3}$$

In this formula, we can simplify the terms, given the high level of symmetry of the equal-mass binary system.

$$\begin{aligned}
m_A &= m_B = m, \\
z_A &= z + d/2 \Rightarrow Z_A = \gamma(z + d/2), \\
z_B &= z - d/2 \Rightarrow Z_B = \gamma(z - d/2), \\
R_A &= \gamma|z + d/2| \rightarrow \gamma(z + d/2), \\
R_B &= \gamma|z - d/2| \rightarrow -\gamma(z - d/2), \\
\zeta_A &= \gamma^2[z + d/2 - v|z + d/2|] \rightarrow \gamma^2(1 - v)(z + d/2), \\
\zeta_B &= \gamma^2[z - d/2 + v|z - d/2|] \rightarrow \gamma^2(1 - v)(z - d/2),
\end{aligned}$$

where we have also assumed that z was in the range $z \in [0, d/2 - \ell]$. Then we find that the

relevant three-metric component reduces to:

$$\begin{aligned}
\gamma_{zz} &= 1 + \frac{2\,m_A\,\zeta_A^2}{R_A^3} + \frac{2\,m_B\,\zeta_B^2}{R_B^3} \\
&= 1 + \frac{2\,m\,\gamma^4(1-v)^2(z+d/2)^2}{\gamma^3(z+d/2)^3} + \frac{2\,m\,\gamma^4(1-v)^2(z-d/2)^2}{-\gamma^3(z-d/2)^3} \\
&= 1 + 2\,m\,\gamma(1-v)^2\left[\frac{1}{(z+d/2)} - \frac{1}{(z-d/2)}\right] \\
&= 1 + \frac{2\,m\,\gamma(1-v)^2 d}{d^2/4 - z^2}.
\end{aligned}$$

So the separation integral is

$$\begin{aligned}
D &= 2\int_{z=0}^{d/2-\ell} \sqrt{\gamma_{zz}}\,dz \\
&= 2\int_{z=0}^{d/2-\ell} \sqrt{1 + \frac{2m\gamma(1-v)^2 d}{d^2/4 - z^2}}\,dz.
\end{aligned}$$

For large separations, we can further approximate the AH radius as $\ell = 2m$, to obtain (3.28). Though the integral has no convenient closed-form solution, we can easily calculate the separation numerically.

Appendix C

Merger Run Parameters

This appendix contains the relevant run parameters and details from the three sets of head-on collisions described in Chapter 5. In each case, I have calculated the approximate proper horizon-to-horizon separation D of the holes on the initial data slice, according to (3.28).

C.1 Coordinate Separation $32m$

Domain size: $16m \times 16m \times 24m$

C.1.1 Initial Data

Newtonian velocity v_N: 0.177.

Hole A center trajectory $z_A(m)$: $16.0 - 0.25t - 2.06 \times 10^{-3}t^2 + 1.83 \times 10^{-5}t^3$.

Proper AH separation D: $30.96m$.

Hole A excision radius R_A: $0.8m$.

Excision merger trigger d_e^*: $2.0m$.

Singularity merger trigger d_s^*: $0.4m$.

C.1.2 First Common AH

The first common AH was found at $t = 44m$.

The common mass was $M_{AH} = 1.97m$.

The equatorial circumference was $C_E = 19.23m$.

The polar circumference was $C_P = 27.58m$.

The AH asymmetry was $\epsilon = -30.3\%$.

C.1.3 Mask Merger

The masks merged at $t = 47.05m$.

The most recent AH was found at $t = 47m$.

The common mass was $M_{AH} = 2.04m$.

The equatorial circumference was $C_E = 24.52m$.

The polar circumference was $C_P = 26.29m$.

The AH asymmetry was $\epsilon = -6.7\%$.

C.2 Coordinate Separation $16m$

Domain Size: $18m \times 18m \times 18m$

C.2.1 Initial Data

Newtonian velocity v_N: 0.25.

Hole A center trajectory $z_A(m)$: $8.0 - 0.23t - 0.01t^2$.

Proper AH separation D: $14.17m$.

Hole A excision radius R_A: $1.0m$.

Excision merger trigger d_e^*: $2.4m$.

Singularity merger trigger d_s^*: $0.4m$.

C.2.2 First Common AH

First common AH found at $t = 16.0m$.

The common mass was $M_{AH} = 1.96m$.

The equatorial circumference was $C_E = 17.78m$.

The polar circumference was $C_P = 28.15m$.

The AH asymmetry was $\epsilon = -36.83\%$

C.2.3 Mask Merger

The masks merged at $t = 17.05m$.

The most recent AH was found at $t = 17.0m$.

The common mass was $M_{AH} = 2.02m$.

The equatorial circumference was $C_E = 20.91m$.

The polar circumference was $C_P = 27.76m$.

The AH asymmetry was $\epsilon = -24.38\%$.

C.2.4 Gauge Merger

The gauge switched to live at $t = 18.75m$.

C.3 Coordinate Separation $12m$

Domain Size: $12m \times 12m \times 12m$.

C.3.1 Initial Data

Newtonian velocity v_N: 0.288.

Hole A trajectory $z_A(t)$: $6.0 - 0.365t - 0.01t^2$.

Proper AH separation D: $9.29m$.

Hole A excision Radius R_A: $1.0m$.

Excision merger trigger d_e^*: $2.4m$.

Singularity merger trigger d_s^*: $0.4m$.

C.3.2 First Common AH

First common AH found at $t = 8.0m$.

The common mass was $M_{AH} = 1.97m$.

The equatorial circumference was $C_E = 16.93m$.

The polar circumference was $C_P = 28.54m$.

The AH asymmetry was $\epsilon = -40.67\%$.

C.3.3 Mask Merger

The masks metered at

The most recent AH was found at $t = 10.0m$.

The common mass was $M_{AH} = 2.11m$.

The equatorial circumference was $C_E = 23.54m$.

The polar circumference was $C_P = 27.89m$.

The AH asymmetry was $\epsilon = -15.59\%$.

References

[1] http://www.cactuscode.org.

[2] http://www.appleswithapples.org.

[3] A Abrahams and R Price. Black-hole collisions from Brill-Lindquist initial data: Predictions of perturbation theory. *Phys. Rev. D*, 53:1972–1976, 1996. gr-qc/9509020.

[4] C Alcock et al. Strange stars. *Astrophys. J.*, 310:261–272, 1986.

[5] M Alcubierre and B Brügmann. Simple excision of a black hole in 3+1 numerical relativity. *Phys. Rev. D*, 63:104006, 2001. gr-qc/0008067.

[6] M Alcubierre et al. 3D grazing collision of two black holes. *Phys. Rev. Lett.*, 87: 271103, 2001. gr-qc/0012079.

[7] A Alekseenko and D Arnold. A new symmetric hyperbolic formulation for the Einstein equations. *Phys. Rev. D*, 68:064013, 2003. gr-qc/0210071.

[8] N Andersson et al. A new numerical approach to the oscillation modes of relativistic stars. *Mon. Not. R. Astron. Soc.*, 274:1039–1048, 1995. gr-qc/9503014.

[9] P Anninos et al. Collision of two black holes. *Phys. Rev. Lett.*, 71:2851–2854, 1993. gr-qc/9309016.

[10] P Anninos et al. Head-on collision of two black holes. *Phys. Rev. D*, 52:2044–2058, 1995.

[11] P Anninos et al. Head-on collision of two black holes: Comparison of different approaches. *Phys. Rev. D*, 52:4462–4480, 1995. gr-qc/9505042.

[12] R Arnowitt et al. The dynamics of general relativity. In Louis Witten, editor, *Gravitation: An Introduction to Current Research*, pages 227–265. New York : Wiley, 1962.

[13] A Ashtekar and B Krishnan. Dynamical horizons and their properties. *Phys. Rev. D*, 68:104030, 2003. gr-qc/0308033.

[14] L Baiotti et al. Three-dimensional relativistic simulations of rotating neutron star collapse to a Kerr black hole. gr-qc/0403029, 2004.

[15] J Baker and M Campanelli. Making use of geometrical invariants in black hole collisions. *Phys. Rev. D*, 62:127501, 2000. gr-qc/0003031.

[16] J Baker et al. Gravitational waves from black hole collisions via an eclectic approach. *Class. Quantum Grav.*, 17:L149–L156, 2000. gr-qc/0003027.

[17] J Baker et al. The Lazarus project: A pragmatic approach to binary black hole evolution. *Phys. Rev. D*, 65:044001, 2002. gr-qc/0104063.

[18] J Baker et al. Modeling gravitational radiation from coalescing binary black holes. *Phys. Rev. D*, 65:124012, 2002. astro-ph/0202469.

[19] T Baumgarte and S Shapiro. Numerical integration of Einstein's field equations. *Phys. Rev. D*, 59:024007, 1998. gr-qc/9810065.

[20] C Beetle and L Burko. A radiation scalar for numerical relativity. *Phys. Rev. Lett.*, 89:271101, 2002. gr-qc/0210019.

[21] E Berti et al. Highly damped quasinormal modes of Kerr black holes. *Phys. Rev. D*, 68:124018, 2003. hep-th/0307013.

[22] E Berti and K Kokkotas. Asymptotic quasinormal modes of Reissner-Nordstrm and Kerr black holes. *Phys. Rev. D*, 68:044027, 2003. hep-th/0303029.

[23] N Bishop et al. Black hole data via a Kerr-Schild approach. *Phys. Rev. D*, 57: 6113–6118, 1998. gr-qc/9711076.

[24] N Bishop et al. Black hole initial data from a nonconformal decomposition. *Phys. Rev. D*, 69:064010, 2004. gr-qc/0310011.

[25] L Blanchet. Gravitational radiation from post-Newtonian sources and inspiralling compact binaries. *Living Reviews in Relativity*, 2002:3, 2002. http://www.livingreviews.org/lrr-2002-3.

[26] C Bona et al. New formalism for numerical relativity. *Phys. Rev. Lett.*, 75:600–603, 1995. gr-qc/9412071.

[27] E Bonning et al. Physics and initial data for multiple black hole spacetimes. *Phys. Rev. D*, 68:044019, 2003. gr-qc/0305071.

[28] J Bowen and J York. Time-asymmetric initial data for black holes and black-hole collisions. *Phys. Rev. D*, 21:2047–2056, 1980.

[29] R Boyer and R Lindquist. Maximal analytic extension of the Kerr metric. *J. Math. Phys.*, 8:265–281, 1967.

[30] S Brandt and B Brügmann. A simple construction of initial data for multiple black holes. *Phys. Rev. Lett.*, 78:3606–3609, 1997. gr-qc/9703066.

[31] S Brandt et al. Grazing collisions of black holes via the excision of singularities. *Phys. Rev. Lett.*, 85:5496–5499, 2000. gr-qc/0009047.

[32] S Brandt and E Seidel. Evolution of distorted rotating black holes. II. Dynamics and analysis. *Phys. Rev. D*, 52:870–886, 1995. gr-qc/9412073.

[33] D Brill and R Lindquist. Interaction energy in geometrostatics. *Phys. Rev.*, 131: 471–476, 1963.

[34] B Brügmann. Adaptive mesh and geodesically sliced Schwarzschild spacetime in 3+1 dimensions. *Phys. Rev. D*, 54:7361–7372, 1996. gr-qc/9608050.

[35] B Brügmann. Binary black hole mergers in 3D numerical relativity. *Int. J. Mod. Phys. D*, 8:85–100, 1999. gr-qc/9708035.

[36] B Brügmann et al. Numerical simulation of orbiting black holes. gr-qc/0312112, 2003.

[37] G Calabrese et al. Novel finite-differencing techniques for numerical relativity: Application to black hole excision. *Class. Quantum Grav.*, 20:L245–L252, 2003. gr-qc/0302072.

[38] G Calabrese and O Sarbach. Detecting ill posed boundary conditions in general relativity. *J. Math. Phys.*, 44:3888–3899, 2003. gr-qc/0303040.

[39] K Camarda and E Seidel. Numerical evolution of dynamic 3D black holes: Extracting waves. *Phys. Rev. D*, 57:R3204, 1998. gr-qc/9709075.

[40] M Campanelli et al. Perturbations of the Kerr spacetime in horizon-penetrating coordinates. *Class. Quantum Grav.*, 18:1543–1554, 2001. gr-qc/0010034.

[41] S M Carroll, W H Press, and E L Turner. The cosmological constant. *Annu. Rev. Astron. Astrophys.*, 30:499–542, 1992.

[42] S Chandrasekhar. *The Mathematical Theory of Black Holes.* OUP, 1983.

[43] LIGO Scientific Collaboration. Analysis of First LIGO Science data for stochastic gravitational waves. gr-qc/0312088, 2003.

[44] LIGO Scientific Collaboration. Analysis of ligo data for gravitational waves from binary neutron stars. gr-qc/0308069, 2003.

[45] LIGO Scientific Collaboration. First upper limits from LIGO on gravitational wave bursts. gr-qc/0312056, 2003.

[46] LIGO Scientific Collaboration. Setting upper limits on the strength of periodic gravitational waves using the first science data from the GEO600 and LIGO detectors. gr-qc/0308050, 2003.

[47] G Cook. Corotating and irrotational binary black holes in quasicircular orbits. *Phys. Rev. D*, 65:084003, 2002. gr-qc/0108076.

[48] G Cook et al. Boosted three-dimensional black-hole evolutions with singularity excision. *Phys. Rev. Lett.*, 80:2512–2516, 1998. gr-qc/9711078.

[49] C T Cunningham et al. Radiation from collapsing relativistic stars. I - Linearized odd-parity radiation. *Astrophys. J.*, 224:643–667, 1978.

[50] C T Cunningham et al. Radiation from collapsing relativistic stars. II - Linearized even-parity radiation. *Astrophys. J.*, 230:870–892, 1979.

[51] C T Cunningham et al. Radiation from collapsing relativistic stars. III - Second order perturbations of collapse with rotation. *Astrophys. J.*, 236:674–692, 1980.

[52] C Cutler and K S Thorne. An overview of gravitational-wave sources. In *Proceedings of GR16, Durban*, 2002. gr-qc/0204090.

[53] S Dain. Initial data for a head-on collision of two Kerr-like black holes with close limit. *Phys. Rev. D*, 64:124002, 2001. gr-qc/0103030.

[54] S Dain. Initial data for two Kerr-like black holes. *Phys. Rev. Lett.*, 87:121102, 2001. gr-qc/0012023.

[55] K Danzmann et al. LISA — laser interferometer space antenna, pre-Phase A report. *Max Planck Institut für Quantum-Optiks*, 1998. Report MPQ 233.

[56] P Diener. A new general purpose event horizon finder for 3D numerical spacetimes. *Class. Quantum Grav.*, 20:4901–4918, 2003. gr-qc/0305039.

[57] O Dreyer et al. Introduction to isolated horizons in numerical relativity. *Phys. Rev. D*, 67:024018, 2003. gr-qc/0206008.

[58] O Dreyer et al. Black-hole spectroscopy: testing general relativity through gravitational-wave observations. *Class. Quantum Grav.*, 21:787–803, 2004. gr-qc/0309007.

[59] D M Eardley et al. Gravitational-wave observations as a tool for testing relativistic gravity. *Phys. Rev. D*, 8:3308–3321, 1973.

[60] E Echeverria. Gravitational-wave measurements of the mass and angular momentum of a black hole. *Phys. Rev. D*, 40:3194–3203, 1989.

[61] A Eddington. A comparison of Whitehead's and Einstein's formulae. *Nature*, 113: 192, 1924.

[62] S Fairhurst et al. In preparation.

[63] G J Feldman and R D Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998.

[64] D Finkelstein. Past-future asymmetry of the gravitational field of a point particle. *Phys. Rev.*, 110:965–967, 1958.

[65] L S Finn. Gravitational waves from solar oscillations: Proposal for a transition-zone test of general relativity. *Class. Quantum Grav.*, 2:381–402, 1985.

[66] L S Finn. Detection, measurement, and gravitational radiation. *Phys. Rev. D*, 46: 5236–5249, 1992.

[67] L S Finn and P J Sutton. Bounding the mass of the graviton using binary pulsar observations. *Phys. Rev. D*, 65:044022, 2002. gr-qc/0109049.

[68] L S Finn and K S Thorne. Gravitational waves from a compact star in a circular, inspiral orbit, in the equatorial plane of a massive, spinning black hole, as observed by LISA. *Phys. Rev. D*, 62:124021, 2000. gr-qc/0007074.

[69] É É Flanagan and S A Hughes. Measuring gravitational waves from binary black hole coalescences. I. Signal to noise for inspiral, merger, and ringdown. *Phys. Rev. D*, 57: 4535–4565, 1998. gr-qc/9701039.

[70] S Frittelli and R Gomez. Ill-posedness in the Einstein equations. *J. Math. Phys.*, 41: 5535–5549, 2000. gr-qc/0006082.

[71] S Frittelli and R Gomez. Boundary conditions for hyperbolic formulations of the Einstein equations. *Class. Quantum Grav.*, 20:2379–2392, 2003. gr-qc/0302032.

[72] S Frittelli and O Reula. First-order symmetric hyperbolic Einstein equations with arbitrary fixed gauge. *Phys. Rev. Lett.*, 76:4667–4670, 1996. gr-qc/9605005.

[73] A Gentle et al. The constraints as evolution equations for numerical relativity. *Class. Quantum Grav.*, 21:83–91, 2004. gr-qc/0307007.

[74] G W Gibbons and K Maeda. Black holes and membranes in higher-dimensional theories with dilaton fields. *Nucl. Phys. B*, 298:741–775, 1988.

[75] C Gundlach and J Martin-Garcia. Symmetric hyperbolicity and consistent boundary conditions for second-order Einstein equations. gr-qc/0403019, 2004.

[76] L Gunnarsen et al. A '3+1' method for finding principal null directions. *Class. Quantum Grav.*, 12:133–140, 1995. gr-qc/9406003.

[77] M G Haehnelt. Hierarchical build-up of galactic bulges and the merging rate of supermassive binary black holes. *Class. Quantum Grav.*, 20:S31–S36, 2003. astro-ph/0307379.

[78] M G Haehnelt. Joint formation of supermassive black holes and galaxies. In *Carnegie Observatories Astrophysics Series, Vol. 1: Coevolution of Black Holes and Galaxies, ed. L. C. Ho*. Cambridge University Press, 2003. astro-ph/0307378.

[79] S Hahn and R Lindquist. The two-body problem in geometrodynamics. *Ann. Phys.*, 29:304–331, 1964.

[80] R Hamming. *Numerical Methods for Scientists and Engineers*. Dover, 1987.

[81] S Hawking and G Ellis. *The Large Scale Structure of Space-Time*. Cambridge University Press, 1973.

[82] M Huq et al. Locating boosted Kerr and Schwarzschild apparent horizons. *Phys. Rev. D*, 66:084024, 2002. gr-qc/0002076.

[83] B Kelly et al. Cure for unstable numerical evolutions of single black holes: Adjusting the standard ADM equations in the spherically symmetric case. *Phys. Rev. D*, 64: 084013, 2001. gr-qc/0103099.

[84] R Kerr and A Schild. Some algebraically degenerate solutions of Einstein's gravitational field equations. In *Proceedings of Symposia in Applied Mathematics Vol XVII*, pages 199–209, 1965.

[85] G Khanna et al. Inspiraling black holes: The close limit. *Phys. Rev. Lett.*, 83:3581–3584, 1999. gr-qc/9905081.

[86] A Khokhlov and I Novikov. Gauge stability of 3+1 formulations of general relativity. *Class. Quantum Grav.*, 19:827–846, 2002. gr-qc/0111023.

[87] L Kidder et al. Extending the lifetime of 3D black hole computations with a new hyperbolic system of evolut ion equations. *Phys. Rev. D*, 64:064017, 2001. gr-qc/0105031.

[88] K Kokkotas and B Schmidt. Quasi-normal modes of stars and black holes. *Living Reviews in Relativity*, 1999:2, 1999. http://www.livingreviews.org/lrr-1999-2.

[89] D Kramer et al. *Exact Solutions of Einstein's Field Equations*. Cambridge University Press, 1980.

[90] E Leaver. An analytic representation for the quasi-normal modes of Kerr black holes. *Proc. R. Soc. London*, 402:285–298, 1985.

[91] P Marronetti et al. Approximate analytical solutions to the initial data problem of black hole binary systems. *Phys. Rev. D*, 62:024017, 2000. gr-qc/0001077.

[92] R Matzner et al. Initial data and coordinates for multiple black hole systems. *Phys. Rev. D*, 59:024015, 1998. gr-qc/9805023.

[93] C Misner. Wormhole initial conditions. *Phys. Rev.*, 118:1110–1111, 1960.

[94] C Misner. The method of images in geometrostatics. *Ann. Phys.*, 24:102–117, 1963.

[95] C Misner, K Thorne, and J Wheeler. *Gravitation*. Freeman, San Francisco, 1973.

[96] C Moreno et al. Kerr-Schild-type initial data for black holes with angular momenta. *Class. Quantum Grav.*, 19:6059–6073, 2002. gr-qc/0205060.

[97] G Nagy et al. Strongly hyperbolic second order Einstein's evolution equations. gr-qc/0402123, 2004.

[98] J Neyman. Outline of a theory of statistical estimation based on the classical theory of probability. *Phil. Trans. R. Soc. Lond. A*, 236:333–380, 1937.

[99] H P Nollert. Quasinormal modes: The characteristic 'sound' of black holes and neutron stars. *Class. Quantum Grav.*, 16:R159–R216, 1999.

[100] H Onozowa. Detailed study of quasinormal frequencies of the Kerr black hole. *Phys. Rev. D*, 55:3593–3602, 1997.

[101] M Pelath et al. Trapped surfaces in prolate collapse in the Gibbons-Penrose construction. *Class. Quantum Grav.*, 15:3917–3934, 1998. gr-qc/9805051.

[102] H Pfeiffer et al. Comparing initial-data sets for binary black holes. *Phys. Rev. D*, 66: 024047, 2002. gr-qc/0203085.

[103] W Press et al. *Numerical Recipes in Fortran, 2nd Edition*. Cambridge University Press, 1992.

[104] W Press and S Teukolsky. Perturbations of a rotating black hole. II. Dynamical stability of the Kerr metric. *Astrophys. J.*, 185:649–673, 1973.

[105] W H Press. Long wave trains of gravitational waves from a vibrating black hole. *Astrophys. J. Lett.*, 170:L105–L108, 1971.

[106] R Price and J Pullin. Colliding black holes: The close limit. *Phys. Rev. D*, 72: 3297–3300, 1994. gr-qc/9402039.

[107] T Regge and J Wheeler. Stability of a Schwarzschild singularity. *Phys. Rev.*, 108: 1063–1069, 1957.

[108] F D Ryan. Accuracy of estimating the multipole moments of a massive body from the gravitational waves of a binary inspiral. *Phys. Rev. D*, 56:1845–1855, 1997.

[109] O Sarbach et al. Close limit evolution of Kerr-Schild type initial data for binary black holes. *Phys. Rev. D*, 65:064026, 2002. gr-qc/0110085.

[110] P D Scharre and C M Will. Testing scalar-tensor gravity using space gravitational-wave interferometers. *Phys. Rev. D*, 65:042002, 2002. gr-qc/0109044.

[111] E Schnetter. *Gauge fixing for the simulation of black hole spacetimes.* PhD thesis, Universität zu Tübingen, 2003.

[112] E Schnetter. Comparing numerically computed spacetimes, 2004. Talk delivered at PSU Numerical Lunch, Spring 2004.

[113] K Schwarzschild. Über das Gravitationsfeld eines Massenpunktes nach der Einsteinschen Theorie. *Sitzber. Deut. Akad. Wiss. Berlin, Kl. Math.-Phys. Tech.*, pages 189–196, 1916.

[114] E Seidel and W Suen. Dynamical evolution of boson stars: Perturbing the ground state. *Phys. Rev. D*, 42:384–403, 1990.

[115] M Shahram and P Milanfar. On the resolvability of multiple sinusoids with nearby frequencies in the presence of noise. IEEE Trans. Signal Process., at press, 2003.

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

[117] D Shoemaker et al. Generic tracking of multiple apparent horizons with level flow. *Phys. Rev. D*, 62:124005, 2000. gr-qc/0004062.

[118] D Shoemaker et al. Moving black holes via singularity excision. *Class. Quantum Grav.*, 20:3729–3744, 2003. gr-qc/0301111.

[119] L Smarr. Gauge conditions, radiation formulae and the two black hole collision. In Larry Smarr, editor, *Sources of Gravitational Radiation: Proceedings of the Battelle Seattle Workshop*, pages 275–291. CUP, 1979.

[120] L L Smarr. *The Structure of General Relativity with a Numerical Illustration: The Collision of Two Black Holes*. PhD thesis, University of Texas, Austin, 1975.

[121] K L Smith. *Dynamic Singularity Excision in Numerical Relativity*. PhD thesis, Pennylvania State University, 2004.

[122] D N Spergel et al. First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Determination of cosmological parameters. *Astrophys.J.Suppl.*, 148: 175–194, 2003. astro-ph/0302209.

[123] U Sperhake et al. Impact of densitized lapse slicings on evolutions of a wobbling black hole. *Phys. Rev. D*, 69:024012, 2004. gr-qc/0307015.

[124] A Stuart and J K Ord. *Kendall's Advanced Theory of Statistics, Vol. 1.* Edward Arnold, London, 1994.

[125] P J Sutton and L S Finn. Bounding the graviton mass with binary pulsar observations. *Class. Quantum Grav.*, 19:1355–1360, 2002. gr-qc/0112018.

[126] J H Taylor and J M Weisberg. A new test of general relativity - gravitational radiation and the binary pulsar PSR 1913+16. *Astrophys. J.*, 253:908–920, 1982.

[127] S Teukolsky. Rotating black holes: Separable wave equations for gravitational and electromagnetic perturbations. *Phys. Rev. Lett.*, 29:1114–1118, 1972.

[128] S Teukolsky. Perturbations of a rotating black hole. I. Fundamental equations for gravitational, electromagnetic, and neutrino-field perturbations. *Astrophys. J.*, 185: 635–647, 1973.

[129] S Teukolsky and W Press. Perturbations of a rotating black hole. III. Interaction of the hole with gravitational and electromagnetic radiation. *Ann. Phys.*, 193:443–461, 1974.

[130] J Thornburg. Coordinates and boundary conditions for the general relativistic initial data problem. *Class. Quantum Grav.*, 4:1119–1131, 1987.

[131] J Thornburg. Finding apparent horizons in numerical relativity. *Phys. Rev. D*, 54: 4899–4918, 1996. gr-qc/9508014.

[132] K Thorne. Multipole expansions of gravitational radiation. *Rev. Mod. Phys.*, 52: 299–339, 1980.

[133] K S Thorne. Disk-accretion onto a black hole. II. Evolution of the hole. *Astrophys. J.*, 191:507–519, 1974.

[134] Vishveshwara C V. Scattering of gravitational radiation by a Schwarzschild black-hole. *Nature*, 227:936–938, 1970.

[135] R Wald. *General Relativity*. University of Chicago Press, 1984.

[136] C M Will. Bounding the mass of the graviton using gravitational-wave observations of inspiralling compact binaries. *Phys. Rev. D*, 57:2061–2068, 1998. gr-qc/9709011.

[137] H Yo et al. Improved numerical stability of stationary black hole evolution calculations. *Phys. Rev. D*, 66:084026, 2002. gr-qc/0209066.

[138] G Yoneda and H Shinkai. Constraint propagation in the family of ADM systems. *Phys. Rev. D*, 63:124019, 2001. gr-qc/0103032.

[139] G Yoneda and H Shinkai. Advantages of a modified ADM formulation: Constraint propagation analysis of the Baumgarte-Shapiro-Shibata-Nakamura system. *Phys. Rev. D*, 66:124003, 2002. gr-qc/0204002.

[140] J York. Gravitational degrees of freedom and the initial-value problem. *Phys. Rev. Lett.*, 26:1656–1658, 1971.

[141] J York. Kinematics and dynamics of general relativity. In Larry Smarr, editor, *Sources of Gravitational Radiation: Proceedings of the Battelle Seattle Workshop*, pages 83–126. CUP, 1979.

[142] J York. Initial data for collisions of black holes and other gravitational miscellany. In C Evans et al., editors, *Frontiers in Numerical Relativity*. CUP, 1989.

[143] F Zerilli. Effective potential for even-parity Regge-Wheeler gravitational perturbation equations. *Phys. Rev. Lett.*, 24:737–738, 1970.

[144] F J Zerilli. Gravitational field of a particle falling in a Schwarzschild geometry analyzed in tensor harmonics. *Phys. Rev. D*, 2:2141–2160, 1970.

Vita

Vita

Bernard J Kelly was born in Dublin, Ireland on the 25th March 1973.

He attended University College, Dublin from September 1991, graduating in September 1995, with a joint First-Class Honours Bachelor's degree in Experimental Physics and Mathematical Physics.

He returned to complete a Master's degree from the same institution, graduating with an exam-based Master's in Mathematical Physics in September 1996.

From September 1996 until the Summer of 1998, he was employed as a Financial Mathematician in Beacon FX Ltd, a currency exchange research company.

He left Beacon FX to commence a doctoral program in Penn State University, in the Fall of 1998.

Bernard Kelly has been a student member of the American Physical Society since 2001.