

Event and Apparent Horizon Finders for 3 + 1 Numerical Relativity

Jonathan Thornburg

School of Mathematics, University of Southampton
Highfield, Southampton SO17 1BJ, U.K.

and

Max Planck Institute for Gravitational Physics
(Albert Einstein Institute)

Am Mühlenberg 1, 14476 Potsdam, Germany

email: J.Thornburg@soton.ac.uk

<http://www.aei.mpg.de/~jthorn>

Living Reviews in Relativity

ISSN 1433-8351

Accepted on 17 April 2007

Published on 1 June 2007

Abstract

Event and apparent horizons are key diagnostics for the presence and properties of black holes. In this article I review numerical algorithms and codes for finding event and apparent horizons in numerically-computed spacetimes, focusing on calculations done using the 3 + 1 ADM formalism. The event horizon of an asymptotically-flat spacetime is the boundary between those events from which a future-pointing null geodesic can reach future null infinity and those events from which no such geodesic exists. The event horizon is a (continuous) null surface in spacetime. The event horizon is defined *nonlocally in time*: it is a global property of the entire spacetime and must be found in a separate post-processing phase *after* all (or at least the nonstationary part) of spacetime has been numerically computed.

There are three basic algorithms for finding event horizons, based on integrating null geodesics *forwards* in time, integrating null geodesics *backwards* in time, and integrating null *surfaces* backwards in time. The last of these is generally the most efficient and accurate.

In contrast to an event horizon, an apparent horizon is defined locally in time in a spacelike slice and depends only on data in that slice, so it can be (and usually is) found *during* the numerical computation of a spacetime. A marginally outer trapped surface (MOTS) in a slice is a smooth closed 2-surface whose future-pointing outgoing null geodesics have zero expansion Θ . An apparent horizon is then defined as a MOTS not contained in any other MOTS. The MOTS condition is a nonlinear elliptic partial differential equation (PDE) for the surface shape, containing the ADM 3-metric, its spatial derivatives, and the extrinsic curvature as coefficients. Most “apparent horizon” finders actually find MOTSs.

There are a large number of apparent horizon finding algorithms, with differing trade-offs between speed, robustness, accuracy, and ease of programming. In axisymmetry, shooting algorithms work well and are fairly easy to program. In slices with no continuous symmetries, spectral integral-iteration algorithms and elliptic-PDE algorithms are fast and accurate, but require good initial guesses to converge. In many cases, Schnetter’s “pretracking” algorithm can greatly improve an elliptic-PDE algorithm’s robustness. Flow algorithms are generally quite slow but can be very robust in their convergence. Minimization methods are slow and relatively inaccurate in the context of a finite differencing simulation, but in a spectral code they can be relatively faster and more robust.

Imprint / Terms of Use

Living Reviews in Relativity are published by the Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, 14476 Potsdam, Germany. ISSN 1433-8351

This review is licensed under a Creative Commons Attribution-Non-Commercial-NoDerivs 2.0 Germany License: <http://creativecommons.org/licenses/by-nc-nd/2.0/de/>

Because a *Living Reviews* article can evolve over time, we recommend to cite the article as follows:

Jonathan Thornburg,
“Event and Apparent Horizon Finders for 3 + 1 Numerical Relativity”,
Living Rev. Relativity, **10**, (2007), 3. [Online Article]: cited [<date>],
<http://www.livingreviews.org/lrr-2007-3>

The date given as <date> then uniquely identifies the version of the article you are referring to.

Article Revisions

Living Reviews supports two different ways to keep its articles up-to-date:

Fast-track revision A fast-track revision provides the author with the opportunity to add short notices of current research results, trends and developments, or important publications to the article. A fast-track revision is refereed by the responsible subject editor. If an article has undergone a fast-track revision, a summary of changes will be listed here.

Major update A major update will include substantial changes and additions and is subject to full external refereeing. It is published with a new publication number.

For detailed documentation of an article’s evolution, please refer always to the history document of the article’s online version at <http://www.livingreviews.org/lrr-2007-3>.

Contents

I	Introduction	5
1	Notation and Terminology	6
2	2-Surface Parameterizations	8
2.1	Level-set-function parameterizations	8
2.2	Strahlkörper parameterizations	8
2.3	Finite-element parameterizations	9
3	Software-Engineering Issues	10
3.1	Software libraries and toolkits	10
3.2	Code reuse and sharing	10
3.3	Using multiple event/apparent horizon finders	11
II	Finding Event Horizons	12
4	Introduction	12
5	Algorithms and Codes for Finding Event Horizons	13
5.1	Integrating null geodesics forwards in time	13
5.1.1	Spherically-symmetric spacetimes	14
5.1.2	Non-spherically-symmetric spacetimes	14
5.2	Integrating null geodesics backwards in time	15
5.3	Integrating null surfaces backwards in time	17
5.3.1	Error bounds: Integrating a pair of surfaces	17
5.3.2	Explicit Strahlkörper surface representation	18
5.3.3	Level-set parameterization	21
6	Summary of Algorithms/Codes for Finding Event Horizons	25
III	Finding Apparent Horizons	26
7	Introduction	26
7.1	Definition	26
7.2	General properties	26
7.3	Trapping, isolated, and dynamical horizons	27
7.4	Description in terms of the 3 + 1 variables	27
7.5	Geometry interpolation	28
7.6	Criteria for assessing algorithms	28
7.7	Local versus global algorithms	29
8	Algorithms and Codes for Finding Apparent Horizons	30
8.1	Zero-finding in spherical symmetry	31
8.2	The shooting algorithm in axisymmetry	31
8.3	Minimization algorithms	33
8.3.1	Spurious local minima	33
8.3.2	Performance	34

8.3.3	Anticipating the formation of a common apparent horizon	35
8.3.4	Summary of minimization algorithms/codes	35
8.4	Spectral integral-iteration algorithms	35
8.4.1	Kemball and Bishop’s modifications of the Nakamura–Kojima–Oohara algorithm	36
8.4.2	Lin and Novak’s variant of the Nakamura–Kojima–Oohara algorithm	37
8.4.3	Summary of spectral integral-iteration algorithms	37
8.5	Elliptic-PDE algorithms	37
8.5.1	Angular coordinates, grid, and boundary conditions	38
8.5.2	Evaluating the expansion Θ	38
8.5.3	Solving the nonlinear elliptic PDE	39
8.5.4	The Jacobian matrix	39
8.5.5	Solving the linear equations	41
8.5.6	Sample results	41
8.5.7	Summary of elliptic-PDE algorithms/codes	41
8.6	Horizon pretracking	44
8.6.1	Constant-expansion surfaces	45
8.6.2	Generalized constant-expansion surfaces	45
8.6.3	Goal functions	46
8.6.4	The pretracking search	46
8.6.5	Sample results	47
8.6.6	Summary of horizon pretracking	47
8.7	Flow algorithms	48
8.7.1	Implicit pseudo-time stepping	50
8.7.2	Varying the flow speed	51
8.7.3	Surface representation and the handling of bifurcations	51
8.7.4	Gundlach’s “fast flow”	52
8.7.5	Summary of flow algorithms/codes	52
9	Summary of Algorithms/Codes for Finding Apparent Horizons	53
10	Acknowledgements	54
A	Solving a Single Nonlinear Algebraic Equation	55
B	The Numerical Integration of Ordinary Differential Equations	56
	References	57

Part I

Introduction

Systems with strong gravitational fields, particularly systems which may contain event horizons and/or apparent horizons, are a major focus of numerical relativity. The usual output of a numerical relativity simulation is some (approximate, discrete) representation of the spacetime geometry (the 4-metric and possibly its derivatives) and any matter fields, but *not* any explicit information about the existence, precise location, or other properties of any event/apparent horizons. To gain this information, we must explicitly *find* the horizons from the numerically-computed spacetime geometry. The subject of this review is numerical algorithms and codes for doing this, focusing on calculations done using the 3 + 1 ADM formalism [14, 163]. Baumgarte and Shapiro [27, Section 6] have also recently reviewed event and apparent horizon finding algorithms. The scope of this review is limited to the *finding* of event/apparent horizons and omits any but the briefest mention of the many uses of this information in gaining physical understanding of numerically-computed spacetimes.

In this review I distinguish between a numerical *algorithm* (an abstract description of a mathematical computation; also often known as a “method” or “scheme”), and a computer *code* (a “horizon finder”, a specific piece of computer software which implements a horizon finding algorithm or algorithms). My main focus is on the algorithms, but I also mention specific codes where they are freely available to other researchers.

In this review I have tried to cover all the major horizon finding algorithms and codes, and to accurately credit the earliest publication of important ideas. However, in a field as large and active as numerical relativity, it is not unlikely that I have overlooked and/or misdescribed some important research. I apologise to anyone whose work I’ve slighted, and I ask readers to help make this a truly “living” review by sending me corrections, updates, and/or pointers to additional work (either their own or others) that I should discuss in future revisions of this review.

The general outline of this review is as follows: In the remainder of Part I, I define notation and terminology (Section 1), discuss how 2-surfaces should be parameterized (Section 2), and outline some of the software-engineering issues that arise in modern numerical relativity codes (Section 3). I then discuss numerical algorithms and codes for finding event horizons (Part II) and apparent horizons (Part III). Finally, in the appendices I briefly outline some of the excellent numerical algorithms/codes available for two standard problems in numerical analysis, the solution of a single nonlinear algebraic equation (Appendix A) and the time integration of a system of ordinary differential equations (Appendix B).

1 Notation and Terminology

Except as noted below, I generally follow the sign and notation conventions of Wald [160]. I assume that all spacetimes are globally hyperbolic, and for event-horizon finding I further assume asymptotic flatness; in this latter context \mathcal{J}^+ is future null infinity. I use the Penrose abstract-index notation, with summation over all repeated indices. 4-indices abc range over all spacetime coordinates $\{x^a\}$, and 3-indices ijk range over the spatial coordinates $\{x^i\}$ in a spacelike slice $t = \text{constant}$. The spacetime coordinates are thus $x^a = (t, x^i)$.

Indices uvw range over generic angular coordinates (θ, ϕ) on S^2 or on a horizon surface. Note that these coordinates are conceptually distinct from the 3-dimensional spatial coordinates x^i . Depending on the context, (θ, ϕ) may or may not have the usual polar-spherical topology. Indices IJK label angular grid points on S^2 or on a horizon surface. These are *2-dimensional* indices: a *single* such index uniquely specifies an angular grid point. δ_{IJ} is the Kronecker delta on the space of these indices or, equivalently, on surface grid points.

For any indices p and q , ∂_p and ∂_{pq} are the coordinate partial derivatives $\partial/\partial x^p$ and $\partial^2/\partial x^p \partial x^q$ respectively; for any coordinates μ and ν , ∂_μ and $\partial_{\mu\nu}$ are the coordinate partial derivatives $\partial/\partial \mu$ and $\partial^2/\partial \mu \partial \nu$ respectively. Δ is the flat-space angular Laplacian operator on S^2 , while Δx refers to a finite-difference grid spacing in some variable x .

g_{ab} is the spacetime 4-metric, and g^{ab} the inverse spacetime 4-metric; these are used to raise and lower 4-indices. Γ_{ab}^c are the 4-Christoffel symbols. \mathcal{L}_v is the Lie derivative along the 4-vector field v^a .

I use the 3+1 ‘‘ADM’’ formalism first introduced by Arnowitt, Deser, and Misner [14]; York [163] gives a general overview of this formalism as it is used in numerical relativity. g_{ij} is the 3-metric defined in a slice, and g^{ij} is the inverse 3-metric; these are used to raise and lower 3-indices. ∇_i is the associated 3-covariant derivative operator, and Γ_{ij}^k are the 3-Christoffel symbols. α and β^i are the 3+1 lapse function and shift vector respectively, so the spacetime line element is

$$ds^2 = g_{ab} dx^a dx^b \tag{1}$$

$$= -(\alpha^2 - \beta_i \beta^i) dt^2 + 2\beta_i dx^i dt + g_{ij} dx^i dx^j. \tag{2}$$

As is common in 3+1 numerical relativity, I follow the sign convention of Misner, Thorne, and Wheeler [112] and York [163] in defining the extrinsic curvature of the slice as $K_{ij} = -\frac{1}{2}\mathcal{L}_n g_{ij} = -\nabla_i n_j$, where n^a is the future-pointing unit normal to the slice. (In contrast, Wald [160] omits the minus signs from this definition.) $K \equiv K_i^i$ is the trace of the extrinsic curvature K_{ij} . m_{ADM} is the ADM mass of an asymptotically flat slice.

I often write a differential operator as $F = F(y, \partial_u y, \partial_{uv} y; g_{ij}, \partial_k g_{ij}, K_{ij})$, where the ‘‘;’’ notation means that F is a (generally nonlinear) algebraic function of the variable y and its 1st and 2nd angular derivatives, and that F also depends on the coefficients g_{ij} , $\partial_k g_{ij}$, and K_{ij} at the apparent horizon position.

There are three common types of spacetimes/slices where numerical event or apparent horizon finding is of interest: spherically-symmetric spacetimes/slices, axisymmetric spacetimes/slices, and spacetimes/slices with no continuous spatial symmetries (no spacelike Killing vectors). I refer to the latter as ‘‘fully generic’’ spacetimes/slices.

In this review I use the abbreviations ‘‘ODE’’ for ordinary differential equation, ‘‘PDE’’ for partial differential equation, ‘‘CE surface’’ for constant-expansion surface, and ‘‘MOTS’’ for marginally outer trapped surface. Names in SMALL CAPITALS refer to horizon finders and other computer software.

When discussing iterative numerical algorithms, it is often convenient to use the concept of an algorithm’s ‘‘radius of convergence’’. Suppose the solution space within which the algorithm is iterating is S . Then given some norm $\|\cdot\|$ on S , the algorithm’s radius of convergence about a solution $s \in S$ is defined as the smallest $r > 0$ such that the algorithm will converge to the

correct solution \mathbf{s} for any initial guess \mathbf{g} with $\|\mathbf{g} - \mathbf{s}\| \leq r$. We only rarely know the exact radius of convergence of an algorithm, but practical experience often provides a rough estimate¹.

¹An algorithm's actual "convergence region" (the set of all initial guesses for which the algorithm converges to the correct solution) may even be fractal in shape. For example, the Julia set is the convergence region of Newton's method on a simple nonlinear algebraic equation.

2 2-Surface Parameterizations

2.1 Level-set-function parameterizations

The most general way to parameterize a 2-surface in a slice is to define a scalar “level-set function” F on some neighborhood of the surface, with the surface itself then being defined as the level set

$$F = 0 \quad \text{on the surface.} \quad (3)$$

Assuming the surface to be orientable, it is conventional to choose F so that $F > 0$ ($F < 0$) outside (inside) the surface. The choice of level-set function for a given surface is non-unique, but in general this is not a problem.

This parameterization is valid for any surface topology including time-dependent topologies. The 2-surface itself can then be found by a standard isosurface-finding algorithm such as the marching-cubes algorithm [105]. (This algorithm is widely used in computer graphics and is implemented in a number of widely-available software libraries.)

2.2 Strahlkörper parameterizations

Most apparent horizon finders, and some event-horizon finders, assume that each connected component of the apparent (event) horizon has S^2 topology. With the exception of toroidal event horizons (discussed in Section 4), this is generally a reasonable assumption.

To parameterize an S^2 surface’s shape, it is common to further assume that we are given (or can compute) some “local coordinate origin” point inside the surface such that the surface’s 3-coordinate shape relative to that point is a “Strahlkörper” (literally “ray body”, or more commonly “star-shaped region”), defined by Minkowski [138, Page 108] as

a region in n -D Euclidean space containing the origin and whose surface, as seen from the origin, exhibits only one point in any direction.

The Strahlkörper assumption is a significant restriction on the horizon’s coordinate shape (and the choice of the local coordinate origin). For example, it rules out the coordinate shape and local coordinate origin illustrated in Figure 1: a horizon with such a coordinate shape about the local coordinate origin could not be found by any horizon finder which assumes a Strahlkörper surface parameterization.




Figure 1: *This figure shows a cross-section of a coordinate shape (the thick curve) which is not a Strahlkörper about the local coordinate origin shown (the large dot). The dashed line shows a ray from the local coordinate origin, which intersects the surface in more than one point.*

For event-horizon finding, algorithms and codes are now available which allow an arbitrary horizon topology with no Strahlkörper assumption (see the discussion in Section 5.3.3 for details).

For apparent horizon finding, the flow algorithms discussed in Section 8.7 theoretically allow any surface shape, although many implementations still make the Strahlkörper assumption. Removing this assumption for other apparent horizon finding algorithms might be a fruitful area for further research.

Given the Strahlkörper assumption, the surface can be explicitly parameterized as

$$r = h(\theta, \phi), \quad (4)$$

where r is the Euclidean distance from the local coordinate origin to a surface point, (θ, ϕ) are generic angular coordinates on the horizon surface (or equivalently on S^2), and the “horizon shape function” $h : S^2 \rightarrow \mathfrak{R}^+$ is a positive real-valued function on the domain of angular coordinates defining the surface shape. Given the choice of local coordinate origin, there is clearly a one-to-one mapping between Strahlkörper 2-surfaces and horizon shape functions.

There are two common ways to discretize a horizon shape function:

Spectral representation

Here we expand the horizon shape function h in an infinite series in some (typically orthonormal) set of basis functions such as spherical harmonics $Y_{\ell m}$ or symmetric trace-free tensors²,

$$h(\theta, \phi) = \sum_{\ell, m} a_{\ell m} Y_{\ell m}(\theta, \phi). \quad (5)$$

This series can then be truncated at some finite order ℓ_{\max} , and the $N_{\text{coeff}} = \ell_{\max}(\ell_{\max} + 1)$ coefficients $\{a_{\ell m}\}$ used to represent (discretely approximate) the horizon shape. For reasonable accuracy, ℓ_{\max} is typically on the order of 8 to 12.

Finite difference representation

Here we choose some finite grid of angular coordinates $\{(\theta_{\kappa}, \phi_{\kappa})\}$, $\kappa = 1, 2, 3, \dots, N_{\text{ang}}$ on S^2 (or equivalently on the surface)³, and represent (discretely approximate) the surface shape by the N_{ang} values

$$\{h(\theta_{\kappa}, \phi_{\kappa})\} \quad \kappa = 1, 2, 3, \dots, N_{\text{ang}}. \quad (6)$$

For reasonable accuracy, N_{ang} is typically on the order of a few thousand.

It is sometimes useful to explicitly construct a level-set function describing a given Strahlkörper. A common choice here is

$$F \equiv r - h(\theta, \phi). \quad (7)$$

2.3 Finite-element parameterizations

Another way to parameterize a 2-surface is via finite elements where the surface is modelled as a triangulated mesh, i.e. as a set of interlinked “vertices” (points in the slice, represented by their spatial coordinates $\{x^i\}$), “edges” (represented by ordered pairs of vertices), and faces. Typically only triangular faces are used (represented as oriented triples of vertices).

A key benefit of this representation is that it allows an arbitrary topology for the surface. However, determining the actual surface topology (e.g. testing for whether or not the surface self-intersects) is somewhat complicated.

This representation is similar to that of Regge calculus [128, 72]⁴, and can similarly be expected to show 2nd order convergence with the surface resolution.

²For convenience of exposition I use spherical harmonics here, but there are no essential differences if other basis sets are used.

³I discuss the choice of this angular grid in more detail in Section 8.5.1.

⁴There has been some controversy over whether, and if so how quickly, Regge calculus converges to the continuum

3 Software-Engineering Issues

Historically, numerical relativists wrote their own codes from scratch. As these became more complex, many researchers changed to working on “group codes” with multiple contributors.

3.1 Software libraries and toolkits

More recently, particularly in work on fully generic spacetimes, where all three spatial dimensions must be treated numerically, there has been a strong trend towards the use of higher-level software libraries and modular “computational toolkits” such as CACTUS [74] (<http://www.cactuscode.org>). These have a substantial learning overhead, but can allow researchers to work much more productively by focusing more on numerical relativity instead of computer-science and software-engineering issues such as parameter-file parsing, parallelization, I/O, etc.

A particularly important area for such software infrastructure is mesh refinement⁵. This is essential to much current numerical-relativity research but is moderately difficult to implement even in only one spatial dimension, and much harder in multiple spatial dimensions. There are now a number of software libraries providing multi-dimensional mesh-refinement infrastructure (sometimes combined with parallelization), such as those listed in Table 1. The CACTUS toolkit can be used in either unigrid or mesh-refinement modes, the latter using a “mesh-refinement driver” such as PAGH or CARPET [134, 131] (<http://www.carpetcode.org>).

In this review I point out event and apparent horizon finders which have been written in particular frameworks and comment on whether they work with mesh refinement.

Program	Reference(s)	Web page
programAMRD/PAMR	none	http://laplace.physics.ubc.ca/Group/Software.html
CARPET	[134]	http://www.carpetcode.org
DAGH/GRACE	[117]	http://www.caip.rutgers.edu/~parashar/DAGH/
PARAMESH	[106]	http://ct.gsfc.nasa.gov/paramesh/Users_manual/amr.html
SAMRAI	[86, 87]	http://www.llnl.gov/CASC/SAMRAI/

Table 1: *This table lists some software toolkits for multi-dimensional mesh refinement. All these toolkits also provide parallelization.*

3.2 Code reuse and sharing

Another important issue is that of code reuse and sharing. It is common for codes to be shared within a research group but relatively uncommon for them to be shared between different (competing) research groups. Even apart from concerns about competitive advantage, without a modular structure and clear documentation it is difficult to reuse another group’s code. The use of a common computational toolkit can greatly simplify such reuse.

If such reuse *can* be accomplished, it becomes much easier for other researchers to build on existing work rather than having to “reinvent the wheel”. As well as the obvious ease of reusing existing code that (hopefully!) already works and has been thoroughly debugged and tested, there is

Einstein equations. (See, for example, the debate between Brewin [40] and Miller [110], and the explicit convergence demonstration of Gentle and Miller [73].) However, Brewin and Gentle [41] seem to have resolved this: Regge calculus does, in fact, converge to the continuum solution, and this convergence is generically 2nd order in the resolution.

⁵See, for example, Choptuik [48], Pretorius [127], Schnetter et al. [134], and Pretorius and Choptuik [126] for general surveys of the uses of, and methods for, mesh refinement in numerical relativity.

another – less obvious – benefit of code sharing: It greatly eases the replication of past work, which is essential as a foundation for new development. That is, without access to another researcher’s code, it can be surprisingly difficult to replicate her results because the success or failure of a numerical algorithm frequently depends on subtle implementation details not described in even the most complete of published papers.

Event and apparent horizon finders are excellent candidates for software reuse: Many numerical-relativity researchers can benefit from using them, and they have a relatively simple interface to an underlying numerical-relativity simulation. Even if a standard computational toolkit is not used, this relatively simple interface makes it fairly easy to port an event or apparent horizon finder to a different code.

Table 2 lists event and apparent horizon finders which are freely available to any researcher.

Toolkit/Program	References	Algorithm Type
CACTUS/EHFINDER <code>cvs -d :pserver:cvs_anon@cvs.aei.mpg.de:/numrelcvs checkout AEIThorns/EHFinder</code>	[60]	Integrate null surfaces backwards in time (Section 5.3)
CACTUS/AHFINDER <code>cvs -d :pserver:cvs_anon@cvs.cactuscode.org:/cactusdevcvs checkout CactusEinstein/AHFinder</code>	[4]	Minimization (Section 8.3) and fast flow (Section 8.7)
CACTUS/AHFINDERDIRECT <code>cvs -d :pserver:cvs_anon@cvs.aei.mpg.de:/numrelcvs checkout AEIThorns/AHFinderDirect</code>	[156]	Elliptic-PDE (Section 8.5)
<code>cvs -d :pserver:cvs_anon@cvs.aei.mpg.de:/numrelcvs -r Erik checkout AEIThorns/AHFinderDirect</code>	[133, 135]	Modified by Schnetter [133, 135] to support pretracking (Section 8.6)
CACTUS/TGRAPPARENTHORIZON2D <code>cvs -d :pserver:cvs_anon@cvs.cactuscode.org:/arrangements checkout TAT_Archive/TGRapparentHorizon2D</code>	[132, 133]	Elliptic-PDE (Section 8.5)
LORENE/AH_FINDER <code>http://www.lorene.obspm.fr/</code>	[104]	spectral integral-iteration algorithm (Section 8.4.2)

Table 2: *This table lists event and apparent horizon finders which are freely available to any researcher, along with the cvs repositories or web pages from which they may be obtained.*

3.3 Using multiple event/apparent horizon finders

It is useful to have multiple event or apparent horizon finders available: Their strengths and weaknesses may complement each other, and the extent of agreement or disagreement between their results can help to estimate the numerical accuracy. For example, Figure 11 shows a comparison between the irreducible masses of apparent horizons in a binary black hole coalescence simulation (Alcubierre et al. [5, Figure 4b]), as computed by two different apparent horizon finders in the CACTUS toolkit, AHFINDER and AHFINDERDIRECT. In this case the two agree to within about 2% for the individual horizons and 0.5% for the common horizon.

Part II

Finding Event Horizons

4 Introduction

The black hole region of an asymptotically-flat spacetime is defined [81, 82] as the set of events from which no future-pointing null geodesic can reach future null infinity (\mathcal{I}^+). The event horizon is defined as the boundary of the black hole region. The event horizon is a null surface in spacetime with (in the words of Hawking and Ellis [82, Page 319]) “a number of nice properties” for studying the causal structure of spacetime.

The event horizon is a *global* property of an entire spacetime and is defined *nonlocally in time*: The event horizon in a slice is defined in terms of (and cannot be computed without knowing) the full *future* development of that slice.

In practice, to find an event horizon in a numerically-computed spacetime, we typically instrument a numerical evolution code to write out data files of the 4-metric. After the evolution (or at least the strong-field region) has reached an approximately-stationary final state, we then compute a numerical approximation to the event horizon in a separate post-processing pass, using the 4-metric data files as inputs.

As a null surface, the event horizon is necessarily continuous. In theory it need not be *anywhere* differentiable⁶, but in practice this behavior rarely occurs⁷: The event horizon is generally smooth except for possibly a finite set of “cusps” where new generators join the surface; the surface normal has a jump discontinuity across each cusp. (The classic example of such a cusp is the “inseam” of the “pair of pants” event horizon illustrated in Figures 4 and 5.)

A black hole is defined as a connected component of the black hole region in a 3+1 slice. The boundary of a black hole (the event horizon) in a slice is a 2-dimensional set of events. Usually this has 2-sphere (S^2) topology. However, numerically simulating rotating dust collapse, Abrahams et al. [1] found that in some cases the event horizon in a slice may be *toroidal* in topology. Lehner et al. [99], and Husa and Winicour [91] have used null (characteristic) algorithms to give a general analysis of the event horizon’s topology in black hole collisions; they find that there is generically a (possibly brief) toroidal phase before the final 2-spherical state is reached. Lehner et al. [100] later calculated movies showing this behavior for several asymmetric black hole collisions.

⁶Chruściel and Galloway [49] showed that if a “cloud of sand” falls into a large black hole, each “sand grain” generates a non-differentiable caustic in the event horizon.

⁷This is a statement about the types of spacetimes usually studied by numerical relativists, not a statement about the mathematical properties of the event horizon itself.

5 Algorithms and Codes for Finding Event Horizons

There are three basic event-horizon finding algorithms:

- Integrate null geodesics *forwards* in time (Section 5.1).
- Integrate null geodesics *backwards* in time (Section 5.2).
- Integrate null *surfaces* backwards in time (Section 5.3).

I describe these in detail in the following.

5.1 Integrating null geodesics forwards in time

The first generation of event-horizon finders were based directly on Hawking’s original definition of an event horizon: an event \mathcal{P} is within the black hole region of spacetime if and only if there is no future-pointing “escape route” null geodesic from \mathcal{P} to \mathcal{J}^+ ; the event horizon is the boundary of the black hole region.

That is, as described by Hughes et al. [88], we numerically integrate the null geodesic equation

$$\frac{d^2 x^a}{d\lambda^2} + \Gamma_{bc}^a \frac{dx^b}{d\lambda} \frac{dx^c}{d\lambda} = 0 \quad (8)$$

(where λ is an affine parameter) forwards in time from a set of starting events and check which events have “escaping” geodesics. For analytical or semi-analytical studies like that of Bishop [31], this is an excellent algorithm.

For numerical work it is straightforward to rewrite the null geodesic equation (8) as a coupled system of two first-order equations, giving the time evolution of photon positions and 3-momenta in terms of the 3 + 1 geometry variables α , β^i , g^{ij} , and their spatial derivatives. These can then be time-integrated by standard numerical algorithms⁸. However, in practice several factors complicate this algorithm.

We typically only know the 3 + 1 geometry variables on a discrete lattice of spacetime grid points, and we only know the 3 + 1 geometry variables themselves, not their spatial derivatives. Therefore we must numerically differentiate the field variables, then interpolate the field variables and their spacetime derivatives to each integration point along each null geodesic. This is straightforward to implement⁹, but the numerical differentiation tends to amplify any numerical noise that may be present in the field variables.

Another complicating factor is that the numerical computations generally only span a finite region of spacetime, so it is not entirely obvious whether or not a given geodesic will eventually reach \mathcal{J}^+ . However, if the final numerically-generated slice contains an apparent horizon, we can use this as an approximation: Any geodesic which is inside this apparent horizon will definitely *not* reach \mathcal{J}^+ , while any other geodesic may be assumed to eventually reach \mathcal{J}^+ if its momentum is directed away from the apparent horizon. If the final slice (or at least its strong-field region) is approximately stationary, the error from this approximation should be small. I discuss this stationarity assumption further in Section 5.3.1.

⁸I briefly review ODE integration algorithms and codes in Appendix B.

⁹In practice the differentiation can usefully be combined with the interpolation; I outline how this can be done in Section 7.5.

5.1.1 Spherically-symmetric spacetimes

In spherical symmetry this algorithm works well and has been used by a number of researchers. For example, Shapiro and Teukolsky [141, 142, 143, 144] used it to study event horizons in a variety of dynamical evolutions of spherically symmetric collapse systems. Figure 2 shows an example of the event and apparent horizons in one of these simulations.




Figure 2: This figure shows part of a simulation of the spherically symmetric collapse of a model stellar core (a $\Gamma = \frac{5}{3}$ polytrope) to a black hole. The event horizon (shown by the dashed line) was computed using the “integrate null geodesics forwards” algorithm described in Section 5.1; solid lines show outgoing null geodesics. The apparent horizon (the boundary of the trapped region, shown shaded) was computed using the zero-finding algorithm discussed in Section 8.1. The dotted lines show the world lines of Lagrangian matter tracers and are labeled by the fraction of baryons interior to them. Figure reprinted with permission from [142]. © 1980 by the American Astronomical Society.

5.1.2 Non-spherically-symmetric spacetimes

In a non-spherically-symmetric spacetime, several factors make this algorithm very inefficient:

- Many trial events must be tried to accurately resolve the event horizon’s shape. (Hughes et al. [88] describe a 2-stage adaptive numerical algorithm for choosing the trial events so as to accurately locate the event horizon as efficiently as possible.)

- At each trial event we must try many different trial-geodesic starting directions to see if any of the geodesics escape to \mathcal{J}^+ (or our numerical approximation to it). Hughes et al. [88] report needing only 48 geodesics per trial event in several nonrotating axisymmetric spacetimes, but about 750 geodesics per trial event in rotating axisymmetric spacetimes, with up to 3000 geodesics per trial event in some regions of the spacetimes.
- Finally, each individual geodesic integration requires many (short) time steps for an accurate integration, particularly in the strong-field region near the event horizon.

Because of these limitations, for non-spherically-symmetric spacetimes the “integrate null geodesics forwards” algorithm has generally been supplanted by the more efficient algorithms I describe in the following.

5.2 Integrating null geodesics backwards in time

It is well-known that future-pointing outgoing null geodesics near the event horizon tend to diverge exponentially in time away from the event horizon. Figure 3 illustrates this behavior for Schwarzschild spacetime, but the behavior is actually quite generic.

Anninos et al. [7] and Libson et al. [103] observed that while this instability is a problem for the “integrate null geodesics forwards in time” algorithm (it forces that algorithm to take quite short time steps when integrating the geodesics), we can turn it to our advantage by integrating the geodesics *backwards* in time: The geodesics will now *converge* on to the horizon¹⁰.

This event-horizon finding algorithm thus integrates a large number of such (future-pointing outgoing) null geodesics backwards in time, starting on the final numerically-generated slice. As the backwards integration proceeds, even geodesics which started far from the event horizon will quickly converge to it. This can be seen, for example, in Figures 2 and 3.

Unfortunately, this convergence property holds only for *outgoing* geodesics. In spherical symmetry the distinction between outgoing and ingoing geodesics is trivial but, as described by Libson et al. [103],

[...] for the general 3D case, when the two tangential directions of the EH are also considered, the situation becomes more complicated. Here normal and tangential are meant in the 3D spatial, not spacetime, sense. Whether or not a trajectory can eventually be “attracted” to the EH, and how long it takes for it to become “attracted,” depends on the photon’s starting direction of motion. We note that even for a photon which is already exactly on the EH at a certain instant, if its velocity at that point has some component tangential to the EH surface as generated by, say, numerical inaccuracy in integration, the photon will move outside of the EH when traced backward in time. For a small tangential velocity, the photon will eventually return to the EH [...] but] the position to which it returns will not be the original position.

This kind of tangential drifting is undesirable not just because it introduces inaccuracy in the location of the EH, but more importantly, because it can lead to spurious dynamics of the “EH” thus found. Neighboring generators may cross, leading to numerically artificial caustic points [...].

Libson et al. [103] also observed:

Another consequence of the second order nature of the geodesic equation is that not just the positions but also the directions must be specified in starting the backward

¹⁰This convergence is only true in a global sense: locally the event horizon has no special geometric properties, and the Riemann tensor components which govern geodesic convergence/divergence may have either sign.

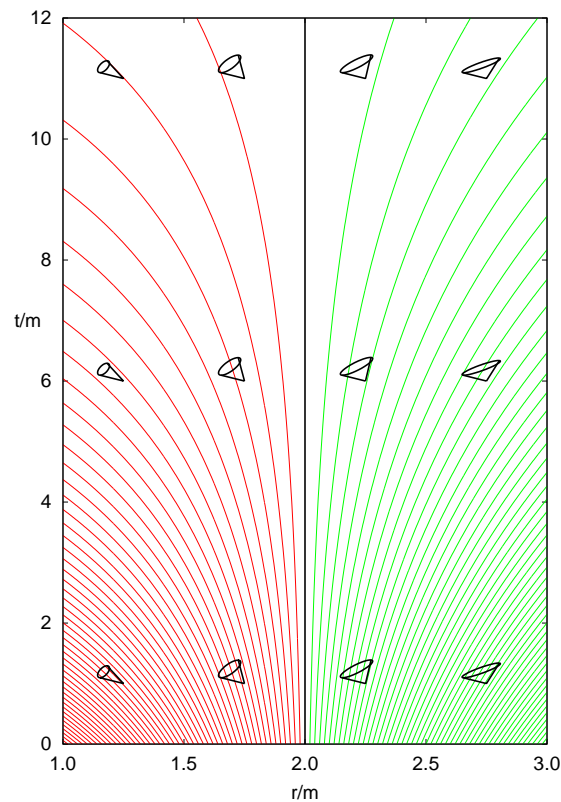


Figure 3: *This figure shows a number of light cones and future-pointing outgoing null geodesics in a neighborhood of the event horizon in Schwarzschild spacetime, plotted in ingoing Eddington–Finkelstein coordinates (t, r) . (These coordinates are defined by the conditions that $t+r$ is an ingoing null coordinate, while r is an areal radial coordinate.) Note that for clarity the horizontal scale is expanded relative to the vertical scale, so the light cones open by more than $\pm 45^\circ$. All the geodesics start out close together near the event horizon; they diverge away from each other exponentially in time (here with an e -folding time of $4m$ near the horizon). Equivalently, they converge towards each other if integrated backwards in time (downwards on the page).*

integration. Neighboring photons must have their starting direction well correlated in order to avoid tangential drifting across one another.

Libson et al. [103] give examples of the numerical difficulties that can result from these difficulties and conclude that this event-horizon finding algorithm

[...] is still quite demanding in finding an accurate history of the EH, although the difficulties are much milder than those arising from the instability of integrating forward in time.

Because of these difficulties, this algorithm has generally been supplanted by the “backwards surface” algorithm I describe next.

5.3 Integrating null surfaces backwards in time

Anninos et al. [7], Libson et al. [103], and Walker [162] introduced the important concept of explicitly (numerically) finding the event horizon as a null *surface* in spacetime. They observed that if we parameterize the event horizon with any level-set function F satisfying the basic level-set definition (3), then the condition for the surface $F = 0$ to be null is just

$$g^{ab} \partial_a F \partial_b F = 0. \quad (9)$$

Applying a 3 + 1 decomposition to this then gives a quadratic equation which can be solved to find the time evolution of the level-set function,

$$\partial_t F = \frac{-g^{ti} \partial_i F + \sqrt{(g^{ti} \partial_i F)^2 - g^{tt} g^{ij} \partial_i F \partial_j F}}{g^{tt}}. \quad (10)$$

Alternatively, assuming the event horizon in each slice to be a Strahlkörper in the manner of Section 2.2, we can define a suitable level-set function F by Equation (7). Substituting this definition into Equation (10) then gives an explicit evolution equation for the horizon shape function,

$$\partial_t h = \frac{-g^{tr} + g^{ru} \partial_u h + \sqrt{(g^{tr} - g^{tu} \partial_u h)^2 - g^{tt} (g^{rr} - 2g^{ru} \partial_u h + g^{uv} \partial_u h \partial_v h)}}{g^{tt}}. \quad (11)$$

Surfaces near the event horizon share the same “attraction” property discussed in Section 5.2 for geodesics near the event horizon. Thus by integrating either surface representation (10) or (11) backwards in time, we can refine an initial guess into a very accurate approximation to the event horizon.

In contrast to the null geodesic equation (8), neither Equation (10) nor Equation (11) contain any derivatives of the 4-metric (or equivalently the 3 + 1 geometry variables). This makes it much easier to integrate these latter equations accurately¹¹. This formulation of the event-horizon finding problem also completely eliminates the tangential-drifting problem discussed in Section 5.2, since the level-set function only parameterizes motion normal to the surface.

5.3.1 Error bounds: Integrating a pair of surfaces

For a practical algorithm, it is useful to integrate a *pair* of trial null surfaces backwards: an “inner-bound” one which starts (and thus always remains) inside the event horizon and an “outer-bound” one which starts (and thus always remains) outside the event horizon. If the final slice contains

¹¹Diener [60] describes how the algorithm can be enhanced to also determine (integrate) individual null generators of the event horizon. This requires interpolating the 4-metric to the generator positions but (still) not taking any derivatives of the 4-metric.

an apparent horizon then any 2-surface inside this can serve as our inner-bound surface. However, choosing an outer-bound surface is more difficult.

It is this desire for a reliable outer bound on the event horizon position that motivates our requirement (Section 4) for the final slice (or at least its strong-field region) to be approximately stationary: In the absence of time-dependent equations of state or external perturbations entering the system, this requirement ensures that, for example, any surface substantially outside the apparent horizon can serve as an outer-bound surface.

Assuming we have an inner- and an outer-bound surface on the final slice, the spacing between these two surfaces after some period of backwards integration then gives an error bound for the computed event horizon position. Equivalently, a necessary (and, if there are no other numerical problems, sufficient) condition for the event-horizon finding algorithm to be accurate is that the backwards integration must have proceeded far enough for the spacing between the two trial surfaces to be “small”. For a reasonable definition of “small”, this typically takes at least $15 m_{\text{ADM}}$ of backwards integration, with $20 m_{\text{ADM}}$ or more providing much higher accuracy.

In some cases it is difficult to obtain a long enough span of numerical data for this backwards integration. For example, in some simulations of binary black hole collisions, the evolution becomes unstable and crashes soon after a common apparent horizon forms. This means that we cannot compute an accurate event horizon for the most interesting region of the spacetime, that which is close to the black-hole merger. There is no good solution to this problem except for the obvious one of developing a stable (or less-unstable) simulation that can be continued for a longer time.

5.3.2 Explicit Strahlkörper surface representation

The initial implementations of the “integrate null surface backwards” algorithm by Anninos et al. [7], Libson et al. [103], and Walker [162] were based on the explicit Strahlkörper surface integration formula (11), further restricted to axisymmetry¹².

For a single black hole the coordinate choice is straightforward. For the two-black-hole case, the authors used topologically cylindrical coordinates (ρ, z, ϕ) , where the two black holes collide along the axisymmetry (z) axis. Based on the symmetry of the problem, they then assumed that the event horizon shape could be written in the form

$$\rho = h(z) \tag{12}$$

in each $t = \text{constant}$ slice.

This spacetime’s event horizon has the now-classic “pair of pants” shape, with a non-differentiable cusp along the “inseam” (the z axis $\rho = 0$) where new generators join the surface. The authors tried two ways of treating this cusp numerically:

- Since the cusp’s location is known *a priori*, it can be treated as a special case in the angular finite differencing, using one-sided numerical derivatives as necessary.
- Alternatively, in 1994 Thorne suggested calculating the *union* of the event horizon and all its null generators (including those which have not yet joined the surface)¹³. This “surface” has a complicated topology (it self-intersects along the cusp), but it is smooth everywhere. This is illustrated by Figure 4, which shows a cross-section of this surface in a single slice, for a head-on binary black hole collision. For comparison, Figure 5 shows a perspective view of part of the event horizon and some of its generators, for a similar head-on binary black hole collision.

¹²Walker [162] mentions an implementation for fully generic slices but only presents results for the axisymmetric case.

¹³See [7, 103, 162].

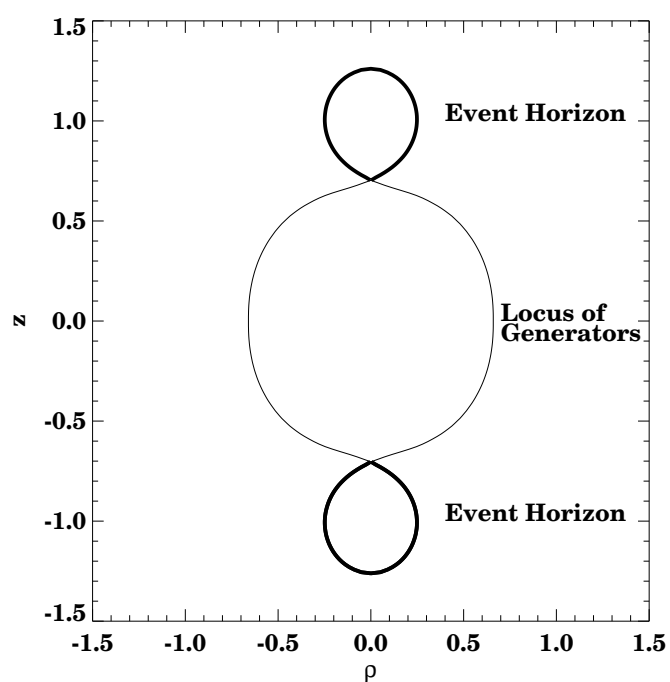


Figure 4: *This figure shows a view of the numerically-computed event horizon in a single slice, together with the locus of the event horizon's generators that have not yet joined the event horizon in this slice, for a head-on binary black hole collision. Notice how the event horizon is non-differentiable at the cusp where the new generators join it. Figure reprinted with permission from [103]. © 1996 by the American Physical Society.*




Figure 5: *This figure shows a perspective view of the numerically-computed event horizon, together with some of its generators, for the head-on binary black hole collision discussed in detail by Matzner et al. [108]. Figure courtesy of Edward Seidel.*

Caveny et al. [44, 46] implemented the “integrate null surfaces backwards” algorithm for fully generic numerically-computed spacetimes using the explicit Strahlkörper surface integration formula (11). To handle moving black holes, they recentered each black hole’s Strahlkörper parameterization (4) on the black hole’s coordinate centroid at each time step.

For single-black-hole test cases (Kerr spacetime in various coordinates), they report typical accuracies of a few percent in determining the event horizon position and area. For binary-black-hole test cases (Kastor–Traschen extremal-charge black hole coalescence with a cosmological constant), they detect black hole coalescence (which appears as a bifurcation in the backwards time integration) by the “necking off” of the surface. Figure 6 shows an example of their results.




Figure 6: *This figure shows the cross-section of the numerically-computed event horizon in each of five different slices, for the head-on collision of two extremal Kastor–Traschen black holes. Figure reprinted with permission from [46]. © 2003 by the American Physical Society.*

5.3.3 Level-set parameterization

Caveny et al. [44, 45] and Diener [60] (independently) implemented the “integrate null surfaces backwards” algorithm for fully generic numerically-computed spacetimes, using the level-set function integration formula (10). Here the level-set function F is initialized on the final slice of the evolution and evolved backwards in time using Equation (10) on (conceptually) the entire numerical grid. (In practice, only a smaller box containing the event horizon need be evolved.)

This surface parameterization has the advantage that the event-horizon topology and (non-)smoothness are completely unconstrained, allowing the numerical study of configurations such as toroidal event horizons (discussed in Section 4). It is also convenient that the level-set function F is defined on the same numerical grid as the spacetime geometry, so that no interpolation is needed for the evolution.

The major problem with this algorithm is that during the backwards evolution, spatial gradients in F tend to steepen into a jump discontinuity at the event horizon¹⁴, eventually causing numerical difficulty.

Caveny et al. [44, 45] deal with this problem by adding an artificial viscosity (i.e. diffusion) term to the level-set function evolution equation, smoothing out the jump discontinuity in F . That is, instead of Equation (10), they actually evolve F via

$$\partial_t F = \text{rhs}_{\text{Eq. (10)}} + \varepsilon^2 \nabla^2 F, \quad (13)$$

where $\text{rhs}_{\text{Eq. (10)}}$ is the right hand side of Equation (10) and ∇^2 is a generic 2nd order linear (elliptic) spatial differential operator, and $\varepsilon > 0$ is a (small) dissipation constant. This scheme works, but the numerical viscosity does seem to lead to significant errors (several percent) in their computed event-horizon positions and areas¹⁵, and even failure to converge to the correct solution for some test cases (e.g. rapidly-spinning Kerr black holes).

Alternatively, Diener [60] developed a technique of periodically reinitializing the level-set function to approximately the signed distance from the event horizon. To do this, he periodically evolves

$$\partial_\lambda F = -\frac{F}{\sqrt{F^2 + 1}} (|\nabla F| - 1) \quad (14)$$

in an unphysical “pseudo-time” λ until an approximate steady state has been achieved. He reports that this works well in most circumstances but can significantly distort the computed event horizon if the $F = 0$ isosurface (the current approximation to the event horizon) is only a few grid points thick in any direction, as typically occurs just around the time of a topology change in the isosurface. He avoids this problem by estimating the minimum thickness of this isosurface and, if it is below a threshold, deferring the reinitialization.

In various tests on analytical data, Diener [60] found this event-horizon finder, EHFINDER, to be robust and highly accurate, typically locating the event horizon to much less than 1% of the 3-dimensional grid spacing. As an example of results obtained with EHFINDER, Figure 7 shows two views of the numerically-computed event horizon for a spiraling binary black hole collision. As another example, Figure 8 shows the numerically-computed event and apparent horizons in the collapse of a rapidly rotating neutron star to a Kerr black hole. (The apparent horizons were computed using the AHFINDERDIRECT code described in Section 8.5.7.)

EHFINDER is implemented as a freely available module (“thorn”) in the CACTUS computational toolkit (see Table 2). It originally worked only with the PUGH unigrid driver, but work is ongoing [61] to enhance it to work with the CARPET mesh-refinement driver [134, 131].

¹⁴Equivalently, Diener [60] observed that the locus of any given nonzero value of the level-set function F is itself a null surface and tends to move (exponentially) closer and closer to the event horizon as the backwards evolution proceeds.

¹⁵They describe how Richardson extrapolation can be used to improve the accuracy of the solutions from $\mathcal{O}(\varepsilon)$ to $\mathcal{O}(\varepsilon^2)$, but it appears that this has not been done for their published results.




Figure 7: *This figure shows two views of the numerically-computed event horizon's cross-section in the orbital plane for a spiraling binary black hole collision. The two orbital-plane dimensions are shown horizontally; time runs upwards. The initial data was constructed to have an approximate helical Killing vector, corresponding to black holes in approximately circular orbits (the $D = 18$ case of Grandclément et al. [78]), with a proper separation of the apparent horizons of 6.9 m . (The growth of the individual event horizons by roughly a factor of 3 in the first part of the evolution is an artifact of the coordinate choice – the black holes are actually in a quasi-equilibrium state.) Figure courtesy of Peter Diener.*




Figure 8: *This figure shows the polar and equatorial radii of the event horizon (solid lines) and apparent horizon (dashed lines) in a numerical simulation of the collapse of a rapidly rotating neutron star to form a Kerr black hole. The dotted line shows the equatorial radius of the stellar surface. These results are from the D4 simulation of Baiotti et al. [21]. Notice how the event horizon grows from zero size while the apparent horizon first appears at a finite size and grows in a spacelike manner. Notice also that both surfaces are flattened due to the rotation. Figure reprinted with permission from [21]. © 2005 by the American Physical Society.*

6 Summary of Algorithms/Codes for Finding Event Horizons

In spherical symmetry the “integrate null geodesics forwards” algorithm (Section 5.1) can be used, although the “integrate null geodesics backwards” and “integrate null surfaces backwards” algorithms (Sections 5.2 and 5.3 respectively) are more efficient.

In non-spherically-symmetric spacetimes the “integrate null surfaces backwards” algorithm (Section 5.3) is clearly the best algorithm known: It is efficient, accurate, and fairly easy to implement. For generic spacetimes, Diener’s event-horizon finder EHFINDER [60] is particularly notable as a freely available implementation of this algorithm as a module (“thorn”) in the widely-used CACTUS computational toolkit (see Table 2).

Part III

Finding Apparent Horizons

7 Introduction

7.1 Definition

Given a (spacelike) $3+1$ slice, a “trapped surface” is defined as a smooth closed 2-surface in the slice whose future-pointing outgoing null geodesics have *negative* expansion Θ . The “trapped region” in the slice is then defined as the union of all trapped surfaces, and the “apparent horizon” is defined as the outer boundary of the trapped region.

While mathematically elegant, this definition is not convenient for numerically finding apparent horizons. Instead, an alternate definition can be used: A MOTS is defined as a smooth (differentiable) closed orientable 2-surface in the slice whose future-pointing outgoing null geodesics have zero expansion Θ .¹⁶ There may be multiple MOTSs in a slice, either nested within each other or intersecting¹⁷. An apparent horizon is then defined as an outermost MOTS in a slice, i.e. a MOTS not contained in any other MOTS. Kriele and Hayward [98] have shown that subject to certain technical conditions, this definition is equivalent to the “outer boundary of the trapped region” one.

Notice that the apparent horizon is defined *locally in time* (it can be computed using only Cauchy data on a spacelike slice), but (because of the requirement that it be closed) *non-locally in space*¹⁸. Hawking and Ellis [82] discuss the general properties of MOTSs and apparent horizons in more detail.

Except for flow algorithms (Section 8.7), all numerical “apparent horizon” finding algorithms and codes actually find MOTSs, and hereinafter I generally follow the common (albeit sloppy) practice in numerical relativity of blurring the distinction between an MOTS and an apparent horizon.

7.2 General properties

Given certain technical assumptions (including energy conditions), the existence of any trapped surface (and hence any apparent horizon) implies that the slice contains a black hole¹⁹. (The converse of this statement is *not* true: An arbitrary (spacelike) slice through a black hole need not contain any apparent horizon²⁰.) However, if an apparent horizon does exist, it necessarily coincides with, or is contained in, an event horizon. In a stationary spacetime the event and apparent horizons coincide.

¹⁶Note that the surface must be smooth *everywhere*. If this condition were not imposed, then MOTSs would lose many of their important properties. For example, even a standard $t = \text{constant}$ slice of Minkowski spacetime contains many non-smooth “MOTSs”: The surface of any closed polyhedron in such a slice satisfies all the other conditions to be an MOTS.

¹⁷Andersson and Metzger [6] have shown that MOTSs can only intersect if they are contained within an outer common MOTS. Szilágyi et al. [151] give a numerical example of such overlapping MOTSs found in a binary black hole collision.

¹⁸As an indication of the importance of the “closed” requirement, Hawking [81] observed that if we consider two spacelike-separated events in Minkowski spacetime, the intersection of their backwards light cones satisfies all the conditions of the MOTS definition, except that it is not closed.

¹⁹The proof is given by Hawking and Ellis [82, Proposition 9.2.8] and by Wald [160, Propositions 12.2.3 and 12.2.4].

²⁰Wald and Iyer [161] proved this by explicitly constructing a family of angularly anisotropic slices in Schwarzschild spacetime which approach arbitrarily close to $r = 0$ yet contain no apparent horizons. However, Schetter and Krishnan [136] have recently studied the behavior of apparent horizons in various anisotropic slices in Schwarzschild and Vaidya spacetimes, finding that the Wald and Iyer behavior seems to be rare.

It is this relation to the event horizon which makes apparent horizons valuable for numerical computation: An apparent horizon provides a useful approximation to the event horizon in a slice, but unlike the event horizon, an apparent horizon is defined locally in time and so can be computed “on the fly” during a numerical evolution.

Given a family of spacelike 3 + 1 slices which foliate part of spacetime, the union of the slices’ apparent horizons (assuming they exist) forms a world-tube²¹. This world-tube is necessarily either null or spacelike. If it is null, this world-tube is slicing-independent (choosing a different family of slices gives the same world-tube, at least so long as each slice still intersects the world-tube in a surface with 2-sphere topology). However, if the world-tube is spacelike, it is *slicing-dependent*: Choosing a different family of slices will in general give a different world-tube²².

7.3 Trapping, isolated, and dynamical horizons

Hayward [83] introduced the important concept of a “trapping horizon” (roughly speaking an apparent horizon world-tube where the expansion becomes negative if the surface is deformed in the inward null direction) along with several useful variants. Ashtekar, Beetle, and Fairhurst [16], and Ashtekar and Krishnan [18] later defined the related concepts of an “isolated horizon”, essentially an apparent horizon world-tube which is null, and a “dynamical horizon”, essentially an apparent horizon world-tube which is spacelike.

These world-tubes obey a variety of local and global conservation laws, and have many applications in analyzing numerically-computed spacetimes. See the references cited above and also Dreyer et al. [63], Ashtekar and Krishnan [19, 20], Gourgoulhon and Jaramillo [76], Booth [36], and Schnetter, Krishnan, and Beyer [137] for further discussions, including applications to numerical relativity.

7.4 Description in terms of the 3 + 1 variables

In terms of the 3 + 1 variables, a MOTS (and thus an apparent horizon) satisfies the condition²³

$$\Theta \equiv \nabla_i s^i + K_{ij} s^i s^j - K = 0, \quad (15)$$

where s^i is the outward-pointing unit 3-vector normal to the surface²⁴. Assuming the Strahlkörper surface parameterization (4), Equation (15) can be rewritten in terms of angular 1st and 2nd derivatives of the horizon shape function h ,

$$\Theta \equiv \Theta(h, \partial_u h, \partial_{uv} h; g_{ij}, \partial_k g_{ij}, K_{ij}) = 0, \quad (16)$$

where Θ is a complicated nonlinear algebraic function of the arguments shown. (Shibata [146] and Thornburg [153, 156] give the $\Theta(h, \partial_u h, \partial_{uv} h)$ function explicitly.)

²¹This world-tube is sometimes called “the apparent horizon”, but this is not standard terminology. In this review I always use the terminology that an MOTS or apparent horizon is a 2-surface contained in a (single) slice.

²²Ashtekar and Galloway [17] have recently proved “a number of physically interesting constraints” on this slicing-dependence.

²³The derivation of this condition is given by (for example) York [164], Gundlach [80, Section IIA], and Baumgarte and Shapiro [27, Section 6.1].

²⁴Notice that in order for the 3-divergence in Equation (15) to be meaningful, s^i (defined only as a field on the MOTS) must be smoothly continued off the surface and extended to a field in some 3-dimensional neighborhood of the surface. The off-surface continuation is non-unique, but it is easy to see that this does not affect the value of Θ on the surface.

7.5 Geometry interpolation

Θ depends on the slice geometry variables g_{ij} , $\partial_k g_{ij}$, and K_{ij} at the horizon position²⁵. In practice these variables are usually only known on the (3-dimensional) numerical grid of the underlying numerical-relativity simulation²⁶, so they must be interpolated to the horizon position and, more generally, to the position of each intermediate-iterate trial shape the apparent horizon finding algorithm tries in the process of (hopefully) converging to the horizon position.

Moreover, usually the underlying simulation gives only g_{ij} and K_{ij} , so g_{ij} must be numerically differentiated to obtain $\partial_k g_{ij}$. As discussed by Thornburg [156, Section 6.1], it is somewhat more efficient to combine the numerical differentiation and interpolation operations, essentially doing the differentiation inside the interpolator²⁷.

Thornburg [156, Section 6.1] argues that for an elliptic-PDE algorithm (Section 8.5), for best convergence of the nonlinear elliptic solver the interpolated geometry variables should be smooth (differentiable) functions of the trial horizon surface position. He argues that that the usual Lagrange polynomial interpolation does not suffice here (in some cases his Newton’s-method iteration failed to converge) because this interpolation gives results which are only piecewise differentiable²⁸. He uses Hermite polynomial interpolation to avoid this problem. Cook and Abrahams [51], and Pfeiffer et al. [124] use bicubic spline interpolation; most other researchers either do not describe their interpolation scheme or use Lagrange polynomial interpolation.

7.6 Criteria for assessing algorithms

Ideally, an apparent horizon finder should have several attributes:

Robust: The algorithm/code should find an (the) apparent horizon in a wide range of numerically-computed slices, without requiring extensive tuning of initial guesses, iteration parameters, etc. This is often relatively easy to achieve for “tracking” the time evolution of an existing apparent horizon (where the most recent previously-found apparent horizon provides an excellent initial guess for the new apparent horizon position) but may be difficult for detecting the appearance of a new (outermost) apparent horizon in an evolution, or for initial-data or other studies where there is no “previous time step”.

Accurate: The algorithm/code should find an (the) apparent horizon to high accuracy and should not report spurious “solutions” (“solutions” which are not actually good approximations to apparent horizons or, at least, to MOTSS).

Efficient: The algorithm/code should be efficient in terms of its memory use and CPU time; in practice CPU time is generally the major constraint. It is often desirable to find apparent horizons at each time step (or, at least, at frequent intervals) during a numerical evolution. For this to be practical the apparent horizon finder must be very fast.

In practice, no apparent horizon finder is perfect in all these dimensions, so trade-offs are inevitable, particularly when ease of programming is considered.

²⁵Or, in the Huq et al. [89, 90] algorithm described in Section 8.5.2, at the local Cartesian grid point positions.

²⁶If the underlying simulation uses spectral methods then the spectral series can be evaluated anywhere, so no actual interpolation need be done, although the term “spectral interpolation” is still often used. See Fornberg [70], Gottlieb and Orszag [75], and Boyd [37] for general discussions of spectral methods, and (for example) Ansorg et al. [12, 11, 10, 13], Bonazzola et al. [35, 33, 34], Grandclément et al. [77], Kidder et al. [95, 96, 97], and Pfeiffer et al. [120, 124, 123, 122] for applications to numerical relativity.

²⁷Conceptually, an interpolator generally works by locally fitting a fitting function (usually a low-degree polynomial) to the data points in a neighborhood of the interpolation point, then evaluating the fitting function at the interpolation point. By evaluating the *derivative* of the fitting function, the $\partial_k g_{ij}$ values can be obtained very cheaply at the same time as the g_{ij} values.

²⁸Thornburg [154, Appendix F] gives a more detailed discussion of the non-smoothness of Lagrange-polynomial interpolation errors.

7.7 Local versus global algorithms

Apparent horizon finding algorithms can usefully be divided into two broad classes:

Local algorithms are those whose convergence is only guaranteed in some (functional) neighborhood of a solution. These algorithms require a “good” initial guess in order to find the apparent horizon. Most apparent horizon finding algorithms are local.

Global algorithms are those which can (in theory, ignoring finite-step-size and other numerical effects) converge to the apparent horizon independent of any initial guess. Flow algorithms (Section 8.7) are the only truly global algorithms. Zero-finding in spherical symmetry (Section 8.1) and shooting in axisymmetry (Section 8.2) are “almost global” algorithms: They require only 1-dimensional searches, which (as discussed in Appendix A) can be programmed to be very robust and efficient. In many cases horizon pretracking (Section 8.6) can semi-automatically find an initial guess for a local algorithm, essentially making the local algorithm behave like an “almost global” one.

One might wonder why local algorithms are ever used, given the apparently superior robustness (guaranteed convergence independent of any initial guess) of global algorithms. There are two basic reasons:

- In practice, local algorithms are much faster than global ones, particularly when “tracking” the time evolution of an existing apparent horizon.
- Due to finite-step-size and other numerical effects, in practice even “global” algorithms may fail to converge to an apparent horizon. (That is, the algorithms may sometimes fail to find an apparent horizon even when one exists in the slice.)

8 Algorithms and Codes for Finding Apparent Horizons

Many researchers have studied the apparent horizon finding problem, and there are a large number of different apparent horizon finding algorithms and codes. Almost all of these require (assume) that any apparent horizon to be found is a Strahlkörper (Section 2) about some local coordinate origin; both finite-difference and spectral parameterizations of the Strahlkörper are common.

For slices with continuous symmetries, special algorithms are sometimes used:

Zero-Finding in Spherical Symmetry (Section 8.1)

In spherical symmetry the apparent horizon equation (16) becomes a 1-dimensional nonlinear algebraic equation, which can be solved by zero-finding.

The Shooting Algorithm in Axisymmetry (Section 8.2)

In axisymmetry the apparent horizon equation (16) becomes a nonlinear 2-point boundary value ODE, which can be solved by a shooting algorithm.

Alternatively, all the algorithms described below for generic slices are also applicable to axisymmetric slices and can take advantage of the axisymmetry to simplify the implementation and boost performance.

For fully generic slices, there are several broad categories of apparent horizon finding algorithms and codes:

Minimization Algorithms (Section 8.3)

These algorithms define a scalar norm on Θ over the space of possible trial surfaces. A general-purpose scalar-function-minimization routine is then used to search trial-surface-shape space for a minimum of this norm (which should give $\Theta = 0$).

Spectral Integral-Iteration Algorithms (Section 8.4)

These algorithms expand the (Strahlkörper) apparent horizon shape function in a spherical-harmonic basis, use the orthogonality of spherical harmonics to write the apparent horizon equation as a set of integral equations for the spectral coefficients, and solve these equations using a functional-iteration algorithm.

Elliptic-PDE Algorithms (Section 8.5)

These algorithms write the apparent horizon equation (16) as a nonlinear elliptic (boundary-value) PDE for the horizon shape and solve this PDE using (typically) standard elliptic-PDE numerical algorithms.

Horizon Pretracking (Section 8.6)

Horizon pretracking solves a slightly more general problem than apparent horizon finding: Roughly speaking, the determination of the smallest $E \geq 0$ such that the equation $\Theta = E$ has a solution, and the determination of that solution. By monitoring the time evolution of E and of the surfaces satisfying this condition, we can determine – *before* it appears – approximately where (in space) and when (in time) a new MOTS *will* appear in a dynamic numerically-evolving spacetime. Horizon pretracking is implemented as a 1-dimensional (binary) search using a slightly-modified elliptic-PDE apparent horizon finding algorithm as a “subroutine”.

Flow Algorithms (Section 8.7)

These algorithms start with a large 2-surface (larger than any possible apparent horizon in the slice) and shrink it inwards using an algorithm which ensures that the surface will stop shrinking when it coincides with the apparent horizon.

I describe the major algorithms and codes in these categories in detail in the following.

8.1 Zero-finding in spherical symmetry

In a spherically symmetric slice, any apparent horizon must also be spherically symmetric, so the apparent horizon equation (16) becomes a 1-dimensional nonlinear algebraic equation $\Theta(h) = 0$ for the horizon radius h . For example, adopting the usual (symmetry-adapted) polar-spherical spatial coordinates $x^i = (r, \theta, \phi)$, we have [154, Equation (B7)]

$$\Theta \equiv \frac{\partial_r g_{\theta\theta}}{g_{\theta\theta} \sqrt{g_{rr}}} - 2 \frac{K_{\theta\theta}}{g_{\theta\theta}} = 0. \quad (17)$$

Given the geometry variables g_{rr} , $g_{\theta\theta}$, $\partial_r g_{\theta\theta}$, and $K_{\theta\theta}$, this equation may be easily and accurately solved using one of the zero-finding algorithms discussed in Appendix A²⁹.

Zero-finding has been used by many researchers, including [141, 142, 143, 144, 119, 47, 139, 9, 154, 155]³⁰. For example, the apparent horizons shown in Figure 2 were obtained using this algorithm. As another example, Figure 9 shows $\Theta(r)$ and h at various times in a (different) spherically symmetric collapse simulation.

8.2 The shooting algorithm in axisymmetry

In an axisymmetric spacetime we can use symmetry-adapted coordinates (θ, ϕ) , so (given the Strahlkörper assumption) without further loss of generality we can write the horizon shape function as $h = h(\theta)$. The apparent horizon equation (16) then becomes a nonlinear 2-point boundary-value ODE for the horizon shape function h [146, Equation (1.1)]

$$\Theta \equiv \Theta(h, \partial_\theta h, \partial_{\theta\theta} h; g_{ij}, \partial_k g_{ij}, K_{ij}) = 0, \quad (18)$$

where $\Theta(h)$ is a nonlinear 2nd order (ordinary) differential operator in h as shown.

Taking the angular coordinate θ to have the usual polar-spherical topology, local smoothness of the apparent horizon gives the boundary conditions

$$\partial_\theta h = 0 \quad \text{at } \theta = 0 \text{ and } \theta = \theta_{\max}, \quad (19)$$

where θ_{\max} is $\pi/2$ if there is “bitant” reflection symmetry across the $z = 0$ plane, or π otherwise.

As well as the more general algorithms described in the following, this may be solved by a shooting algorithm³¹:

1. Guess the value of h at one endpoint, say $h(\theta=0) \equiv h_*$.
2. Use this guessed value of $h(\theta=0)$ together with the boundary condition (19) there as initial data to integrate (“shoot”) the ODE (18) from $\theta=0$ to the other endpoint $\theta=\theta_{\max}$. This can be done easily and efficiently using one of the ODE codes described in Appendix B.
3. If the numerically computed solution satisfies the other boundary condition (19) at $\theta=\theta_{\max}$ to within some tolerance, then the just-computed $h(\theta)$ describes the (an) apparent horizon, and the algorithm is finished.
4. Otherwise, adjust the guessed value $h(\theta=0) \equiv h_*$ and try again. Because there is only a single parameter (h_*) to be adjusted, this can be done using one of the 1-dimensional zero-finding algorithms discussed in Appendix A.

²⁹Note that $\partial_r g_{\theta\theta}$ is a known coefficient field here, not an unknown; if necessary, it can be obtained by numerically differentiating $g_{\theta\theta}$. Therefore, despite the appearance of the derivative, Equation (17) is still an *algebraic* equation for the horizon radius h , not a differential equation.

³⁰See also the work of Bizoń, Malec, and Ó Murchadha [32] for an interesting analytical study giving necessary and sufficient conditions for apparent horizons to form in non-vacuum spherically symmetric spacetimes.

³¹Ascher, Mattheij, and Russel [15, Chapter 4] give a more detailed discussion of shooting methods.

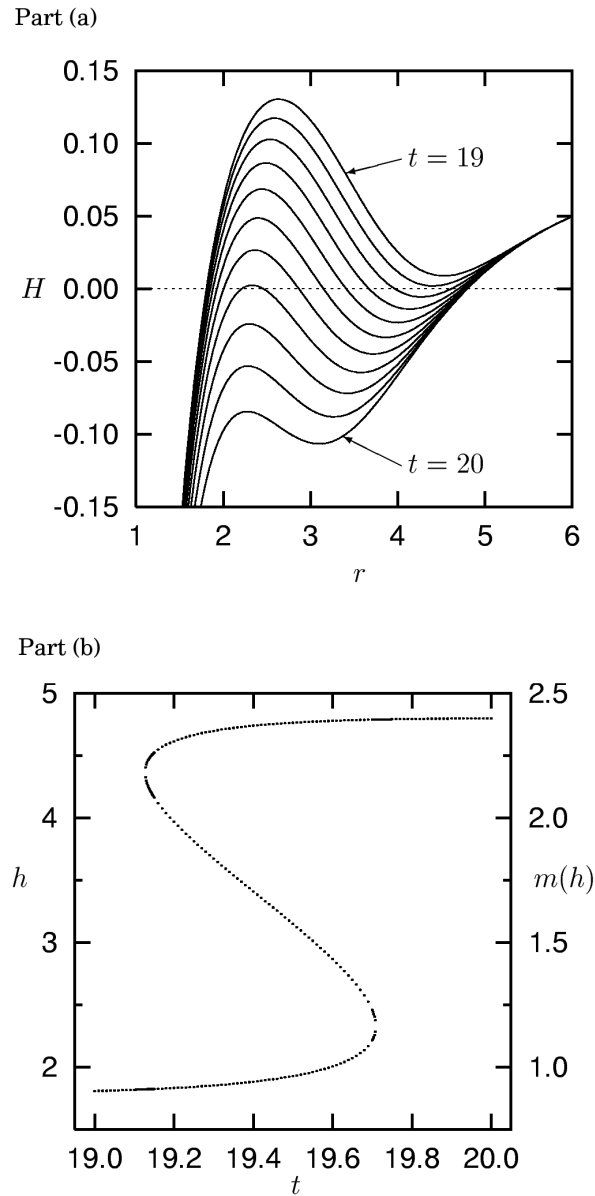


Figure 9: This figure shows the apparent horizons (actually MOTSs) for a spherically symmetric numerical evolution of a black hole accreting a narrow shell of scalar field, the 800.pqw1 evolution of Thornburg [155]. Part (a) of this figure shows $\Theta(r)$ (here labelled H) for a set of equally-spaced times between $t=19$ and $t=20$, while Part (b) shows the corresponding MOTS radius $h(t)$ and the Misner-Sharp [111], [112, Box 23.1] mass $m(h)$ internal to each MOTS. Notice how two new MOTSs appear when the local minimum in $\Theta(r)$ touches the $\Theta=0$ line, and two existing MOTS disappear when the local maximum in $\Theta(r)$ touches the $\Theta=0$ line.

This algorithm is fairly efficient and easy to program. By trying a sufficiently wide range of initial guesses h_* this algorithm can give a high degree of confidence that all apparent horizons have been located, although this, of course, increases the cost.

Shooting algorithms of this type have been used by many researchers, for example [159, 66, 2, 29, 30, 145, 3, 4].

8.3 Minimization algorithms

This family of algorithms defines a scalar norm $\|\cdot\|$ on the expansion Θ over the space of possible trial surfaces, typically the mean-squared norm

$$\|\Theta\| \equiv \int \Theta^2 d\Omega, \quad (20)$$

where the integral is over all solid angles on a trial surface.

Assuming the horizon surface to be a Strahlkörper and adopting the spectral representation (5) for the horizon surface, we can view the norm (20) as being defined on the space of spectral coefficients $\{a_{\ell m}\}$.

This norm clearly has a global minimum $\|\Theta\| = 0$ for each solution of the apparent horizon equation (16). To find the apparent horizon we numerically search the spectral-coefficient space for this (a) minimum, using a general-purpose “function-minimization” algorithm (code) such as Powell’s algorithm³².

Evaluating the norm (20) requires a numerical integration over the horizon surface: We choose some grid of N_{ang} points on the surface, interpolate the slice geometry fields (g_{ij} , $\partial_k g_{ij}$, and K_{ij}) to this grid (see Section 7.5), and use numerical quadrature to approximate the integral. In practice this must be done for many different trial surface shapes (see Section 8.3.2), so it is important that it be as efficient as possible. Anninos et al. [8] and Baumgarte et al. [26] discuss various ways to optimize and/or parallelize this calculation.

Unfortunately, minimization algorithms have two serious disadvantages for apparent horizon finding: They can be susceptible to spurious local minima, and they’re very slow at high angular resolutions. However, for the (fairly common) case where we want to find a common apparent horizon as soon as it appears in a binary black-hole (or neutron-star) simulation, minimization algorithms do have a useful ability to “anticipate” the formation of the common apparent horizon, in a manner similar to the pretracking algorithms discussed in Section 8.6. I discuss the properties of minimization algorithms further in the following.

8.3.1 Spurious local minima

While the norm (20) clearly has a single *global* minimum $\|\Theta\| = 0$ for each MOTS $\Theta = 0$, it typically also has a large number of other *local* minima with $\Theta \neq 0$, which are “spurious” in the sense that they do not correspond (even approximately) to MOTSs³³. Unfortunately, general-purpose “function-minimization” routines only locate local minima and, thus, may easily converge to one of the spurious $\Theta \neq 0$ minima.

What this problem means in practice is that a minimization algorithm needs quite a good (accurate) initial guess for the horizon shape in order to ensure that the algorithm converges to the true global minimum $\Theta = 0$ rather than to one of the spurious $\Theta \neq 0$ local minima.

³²See, for example, Dennis and Schnabel [59], or Brent [39] for general surveys of general-purposes function-minimization algorithms and codes.

³³There is a simple heuristic argument (see, for example, Press et al. [125, Section 9.6]) that at least some spurious local minima should be expected. We are trying to solve a system of N_{ang} nonlinear equations, $\Theta_I = 0$ (one equation for each horizon-surface grid point). Equivalently, we are trying to find the intersection of the N_{ang} codimension-one hypersurfaces $\Theta_I = 0$ in surface-shape space. The problem is that anywhere two or more of these hypersurfaces approach close to each other, but do not actually intersect, there is a spurious local minimum in $\|\Theta\|$.

To view this problem from a different perspective, once the function-minimization algorithm does converge, we must somehow determine whether the “solution” found is the true one, $\Theta = 0$, or a spurious one, $\Theta \neq 0$. Due to numerical errors in the geometry interpolation and the evaluation of the integral (20), $\|\Theta\|$ will almost never evaluate to *exactly* zero; rather, we must set a tolerance level for how large $\|\Theta\|$ may be. Unfortunately, in practice it is hard to choose this tolerance: If it is too small, the genuine solution may be falsely rejected, while if it is too large, we may accept a spurious solution (which may be very different from any of the true solutions).

Anninos et al. [8] and Baumgarte et al. [26] suggest screening out spurious solutions by repeating the algorithm with varying resolutions of the horizon-surface grid and checking that $\|\Theta\|$ shows the proper convergence towards zero. This seems like a good strategy, but it is tricky to automate and, again, it may be difficult to choose the necessary error tolerances in advance.

When the underlying simulation is a spectral one, Pfeiffer et al. [124, 121] report that in practice, spurious solutions can be avoided by a combination of two factors:

- The underlying spectral solution can inherently be “interpolated” (evaluated at arbitrary positions) to very high accuracy.
- Pfeiffer et al. use a large number of quadrature points (typically an order of magnitude larger than the number of coefficients in the expansion (5)) in numerically evaluating the integral (20).

8.3.2 Performance

For convenience of exposition, suppose the spectral representation (5) of the horizon-shape function h uses spherical harmonics $Y_{\ell m}$. (Symmetric trace-free tensors or other basis sets do not change the argument in any important way.) If we keep harmonics up to some maximum degree ℓ_{\max} , the number of coefficients is then $N_{\text{coeff}} = (\ell_{\max} + 1)^2$. ℓ_{\max} is set by the desired accuracy (angular resolution) of the algorithm and is typically on the order of 6 to 12.

To find a minimum in an N_{coeff} -dimensional space (here the space of surface-shape coefficients $\{a_{\ell m}\}$), a general-purpose function-minimization algorithm typically needs on the order of $5N_{\text{coeff}}^2$ to $10N_{\text{coeff}}^2$ iterations³⁴. Thus the number of iterations grows as ℓ_{\max}^4 .

Each iteration requires an evaluation of the norm (20) for some trial set of surface-shape coefficients $\{a_{\ell m}\}$, which requires $\mathcal{O}(N_{\text{coeff}}) = \mathcal{O}(\ell_{\max}^2)$ work to compute the surface positions, together with $\mathcal{O}(N_{\text{ang}})$ work to interpolate the geometry fields to the surface points and compute the numerical quadrature of the integral (20).

Thus the total work for a single horizon finding is $\mathcal{O}(\ell_{\max}^6 + N_{\text{ang}} \ell_{\max}^4)$. Fortunately, the accuracy with which the horizon is found generally improves rapidly with ℓ_{\max} , sometimes even exponentially³⁵. Thus, relatively modest values of ℓ_{\max} (typically in the range 8–12) generally

³⁴A simple counting argument suffices to show that any general-purpose function-minimization algorithm in n dimensions must involve at least $\mathcal{O}(n^2)$ function evaluations (see, for example, Press et al. [125, Section 10.6]): Suppose the function to be minimized is $f : \mathbb{R}^n \rightarrow \mathbb{R}$, and suppose f has a local minimum near some point $x_0 \in \mathbb{R}^n$. Taylor-expanding f in a neighborhood of x_0 gives $f(x) = f(x_0) + \mathbf{a}^T(x-x_0) + (x-x_0)^T \mathbf{B}(x-x_0) + \mathcal{O}(\|x-x_0\|^3)$, where $\mathbf{a} \in \mathbb{R}^n$, $\mathbf{B} \in \mathbb{R}^{n \times n}$ is symmetric, and \mathbf{v}^T denotes the transpose of the column vector $\mathbf{v} \in \mathbb{R}^n$.

Neglecting the higher order terms (i.e. approximating f as a quadratic form in x in a neighborhood of x_0), and ignoring $f(x_0)$ (which does not affect the position of the minimum), there are a total of $N = n + \frac{1}{2}n(n+1)$ coefficients in this expression. Changing any of these coefficients may change the position of the minimum, and at each function evaluation the algorithm “learns” only a single number (the value of f at the selected evaluation point), so the algorithm must make at least $N = \mathcal{O}(n^2)$ function evaluations to (implicitly) determine all the coefficients.

Actual functions are not exact quadratic forms, so in practice there are additional $\mathcal{O}(1)$ multiplicative factors in the number of function evaluations. Minimization algorithms may also make additional performance and/or space-versus-time trade-offs to improve numerical robustness or to avoid explicitly manipulating $n \times n$ Jacobian matrices.

³⁵In the context of an underlying simulation with spectral accuracy, Pfeiffer [122, 121] reports exponential convergence of the horizon finding accuracy with ℓ_{\max} .

suffice for adequate accuracy. Even so, minimization horizon finders tend to be slower than other methods, particularly if high accuracy is required (large ℓ_{\max} and N_{ang}). The one exception is in axisymmetry, where only spherical harmonics $Y_{\ell m}$ with $m=0$ need be considered. In this case minimization algorithms are much faster, though probably still slower than shooting or elliptic-PDE algorithms.

8.3.3 Anticipating the formation of a common apparent horizon

Consider the case where we want to find a common apparent horizon as soon as it appears in a binary black-hole (or neutron-star) simulation. In Section 8.6 I discuss “horizon pretracking” algorithms which can determine – before it appears – approximately where (in space) and when (in time) the common apparent horizon will appear.

Minimization algorithms can provide a similar functionality: Before the common apparent horizon forms, trying to find it via a minimization algorithm will (hopefully) find the (a) surface which minimizes the error norm $\|\Theta\|$ (defined by Equation (20)). This surface can be viewed as the current slice’s closest approximation to a common apparent horizon, and as the evolution proceeds, it should converge to the actual common apparent horizon.

However, it is not clear whether minimization algorithms used in this way suffer from the problems discussed in Section 8.6.2. In particular, it is not clear whether, in a realistic binary-coalescence simulation, the minimum- $\|\Theta\|$ surfaces would remain smooth enough to be represented accurately with a reasonable ℓ_{\max} .

8.3.4 Summary of minimization algorithms/codes

Minimization algorithms are fairly easy to program and have been used by many researchers, for example [43, 69, 102, 8, 26, 4]. However, at least when the underlying simulation uses finite differencing, minimization algorithms are susceptible to spurious local minima, have relatively poor accuracy, and tend to be quite slow. I believe that the other algorithms discussed in the following sections are generally preferable. If the underlying simulation uses spectral methods, then minimization algorithms may be (relatively) somewhat more efficient and robust.

Alcubierre’s apparent horizon finder AHFINDER [4] includes a minimization algorithm based on the work of Anninos et al. [8]³⁶. It is implemented as a freely available module (“thorn”) in the CACTUS computational toolkit (see Table 2).

8.4 Spectral integral-iteration algorithms

Nakamura, Kojima, and Oohara [113] developed a functional-iteration spectral algorithm for solving the apparent horizon equation (16).

This algorithm begins by choosing the usual polar-spherical topology for the angular coordinates (θ, ϕ) , and rewriting the apparent horizon equation (16) in the form

$$\Delta h \equiv \partial_{\theta\theta} h + \frac{\partial_{\theta} h}{\tan \theta} + \frac{\partial_{\phi\phi} h}{\sin^2 \theta} = G(\partial_{\theta\phi} h, \partial_{\phi\phi} h, \partial_{\theta} h, \partial_{\phi} h; g_{ij}, K_{ij}, \Gamma_{ij}^k), \quad (21)$$

where Δ is the flat-space angular Laplacian operator, and G is a complicated nonlinear algebraic function of the arguments shown, which remains regular even at $\theta=0$ and $\theta=\pi$.

Next we expand h in spherical harmonics (5). Because the left hand side of Equation (21) is just the flat-space angular Laplacian of h , which has the $Y_{\ell m}$ as orthogonal eigenfunctions, multiplying

³⁶AHFINDER also includes a “fast flow” algorithm (Section 8.7).

both sides of Equation (21) by $Y_{\ell m}^*$ (the complex conjugate of $Y_{\ell m}$) and integrating over all solid angles gives

$$a_{\ell m} = -\frac{1}{\ell(\ell+1)} \int Y_{\ell m}^* G d\Omega \quad (22)$$

for each ℓ and m except $\ell = m = 0$.

Based on this, Nakamura, Kojima, and Oohara [113] proposed the following functional-iteration algorithm for solving Equation (21):

1. Start with some (initial-guess) set of horizon-shape coefficients $\{a_{\ell m}\}$. These determine a surface shape via Equation (5).
2. Interpolate the geometry variables to this surface shape (see Section 7.5).
3. For each ℓ and m except $\ell = m = 0$, evaluate the integral (22) by numerical quadrature to obtain a next-iteration coefficient $a_{\ell m}$.
4. Determine a next-iteration coefficient a_{00} by numerically solving (finding a root of) the equation

$$\int Y_{00}^* G d\Omega = 0. \quad (23)$$

This can be done using any of the 1-dimensional zero-finding algorithms discussed in Appendix A.

5. Iterate until all the coefficients $\{a_{\ell m}\}$ converge.

Gundlach [80] observed that the subtraction and inversion of the flat-space angular Laplacian operator in this algorithm is actually a standard technique for solving nonlinear elliptic PDEs by spectral methods. I discuss this observation and its implications further in Section 8.7.4.

Nakamura, Kojima, and Oohara [113] report that their algorithm works well, but Nakao (cited as personal communication in [146]) has argued that it tends to become inefficient (and possibly inaccurate) for large ℓ (high angular resolution) because the $Y_{\ell m}$ fail to be numerically orthogonal due to the finite resolution of the numerical grid. I know of no other published work addressing Nakao's criticism.

8.4.1 **Kemball and Bishop's modifications of the Nakamura–Kojima–Oohara algorithm**

Kemball and Bishop [93] investigated the behavior of the Nakamura–Kojima–Oohara algorithm and found that its (only) major weakness seems to be that the a_{00} -update equation (23) “may have multiple roots or minima even in the presence of a single marginally outer trapped surface, and all should be tried for convergence”.

Kemball and Bishop [93] suggested and tested several modifications to improve the algorithm's convergence behavior. They verified that (either in its original form or with their modifications) the algorithm's rate of convergence (number of iterations to reach a given error level) is roughly independent of the degree ℓ_{\max} of spherical-harmonic expansion used. They also give an analysis that the algorithm's cost is $\mathcal{O}(\ell_{\max}^4)$, and its accuracy $\varepsilon = \mathcal{O}(1/\ell_{\max})$, i.e. the cost is $\mathcal{O}(1/\varepsilon^4)$. This accuracy is surprisingly low: Exponential convergence with ℓ_{\max} is typical of spectral algorithms and would be expected here. I do not know of any published work which addresses this discrepancy.

8.4.2 Lin and Novak's variant of the Nakamura–Kojima–Oohara algorithm

Lin and Novak [104] have developed a variant of the Nakamura–Kojima–Oohara algorithm which avoids the need for a separate search for a_{00} at each iteration: Write the apparent horizon equation (16) in the form

$$\Delta h - 2h = \lambda\Theta + \Delta h - 2h, \quad (24)$$

where Δ is again the flat-space angular Laplacian operator and where λ is a nonzero scalar function on the horizon surface. Then choose λ as

$$\lambda = \left(\frac{\det g_{ij}}{\det f_{ij}} \right)^{1/3} [g^{mn}(\bar{\nabla}_m F)(\bar{\nabla}_n F)]^{1/2} h^2, \quad (25)$$

where f_{ij} is the flat metric of polar spherical coordinates, $\bar{\nabla}$ is the associated 3-covariant derivative operator, and F is the level set function (7).

Lin and Novak [104] showed that *all* the spherical-harmonic coefficients $a_{\ell m}$ (including a_{00}) can then be found by iteratively solving the equation

$$a_{\ell m} = -\frac{1}{\ell(\ell+1) + 2} \int Y_{\ell m}^* \lambda (\Theta + \Delta h - 2h) d\Omega. \quad (26)$$

Lin and Novak [104] find that this algorithm gives robust convergence and is quite fast, particularly at modest accuracy levels. For example, running on a 2 GHz processor, their implementation takes 3.1, 5.8, 17, 88, and 313 seconds to find the apparent horizon in a test slice to a relative error (measured in the horizon area) of 9×10^{-4} , 5×10^{-5} , 6×10^{-7} , 9×10^{-10} , and 3×10^{-12} respectively³⁷. This implementation is freely available as part of the LORENE toolkit for spectral computations in numerical relativity (see Table 2).

8.4.3 Summary of spectral integral-iteration algorithms

Despite what appears to be fairly good numerical behavior and reasonable ease of implementation, the original Nakamura–Kojima–Oohara algorithm has not been widely used apart from later work by its original developers (see, for example, [115, 114]). Kemball and Bishop [93] have proposed and tested several modifications to the basic Nakamura–Kojima–Oohara algorithm. Lin and Novak [104] have developed a variant of the Nakamura–Kojima–Oohara algorithm which avoids the need for a separate search for the a_{00} coefficient at each iteration. Their implementation of this variant is freely available as part of the LORENE toolkit for spectral computations in numerical relativity (see Table 2).

8.5 Elliptic-PDE algorithms

The basic concept of elliptic-PDE algorithms is simple: We view the apparent horizon equation (16) as a nonlinear elliptic PDE for the horizon shape function h on the angular-coordinate space and solve this equation by standard finite-differencing techniques³⁸, generally using Newton's method to solve the resulting set of nonlinear algebraic (finite-difference) equations. Algorithms of this type have been widely used both in axisymmetry and in fully generic slices.

³⁷For comparison, the elliptic-PDE AHFINDERDIRECT horizon finder (discussed in Section 8.5.6), running on a roughly similar processor, takes about 1.8 seconds to find the apparent horizon in a similar test slice to a relative error of 4×10^{-4} .

³⁸In theory this equation could also be solved by a spectral method on S^2 , using spectral differentiation to evaluate the angular derivatives. (See [70, 75, 37, 12, 11, 10, 13, 35, 33, 34, 77, 95, 96, 97, 120, 124, 123, 122] for further discussion of spectral methods.) This should yield a highly efficient apparent horizon finder. However, I know of no published work taking this approach.

8.5.1 Angular coordinates, grid, and boundary conditions

In more detail, elliptic-PDE algorithms assume that the horizon is a Strahlkörper about some local coordinate origin, and choose an angular coordinate system and a finite-difference grid of N_{ang} points on S^2 in the manner discussed in Section 2.2.

The most common choices are the usual polar-spherical coordinates (θ, ϕ) and a uniform “latitude/longitude” grid in these coordinates. Since these coordinates are “unwrapped” relative to the actual S^2 trial-horizon-surface topology, the horizon shape function h satisfies periodic boundary conditions across the artificial grid boundary at $\phi = 0$ and $\phi = 2\pi$. The north and south poles $\theta = 0$ and $\theta = \pi$ are trickier, but Huq et al. [89, 90], Shibata and Uryū [147], and Schnetter [132, 133]³⁹ “reaching across the pole” boundary conditions for these artificial grid boundaries.

Alternatively, Thornburg [156] avoids the z axis coordinate singularity of polar-spherical coordinates by using an “inflated-cube” system of six angular patches to cover S^2 . Here each patch’s nominal grid is surrounded by a “ghost zone” of additional grid points where h is determined by interpolation from the neighboring patches. The interpatch interpolation thus serves to tie the patches together, enforcing the continuity and differentiability of h across patch boundaries. Thornburg reports that this scheme works well but was quite complicated to program.

Overall, the latitude/longitude grid seems to be the superior choice: it works well, is simple to program, and eases interoperability with other software.

8.5.2 Evaluating the expansion Θ

The next step in the algorithm is to evaluate the expansion Θ given by Equation (16) on the angular grid given a trial horizon surface shape function h on this same grid (6).

Most researchers compute Θ via 2-dimensional angular finite differencing of Equation (16) on the trial horizon surface. 2nd order angular finite differencing is most common, but Thornburg [156] uses 4th order angular finite differencing for increased accuracy.

With a (θ, ϕ) latitude/longitude grid the $\Theta(h, \partial_u h, \partial_{uv} h)$ function in Equation (16) is singular on the z axis (at the north and south poles $\theta = 0$ and $\theta = \pi$) but can be regularized by applying L’Hôpital’s rule. Schnetter [132, 133] observes that using a *Cartesian* basis for all tensors greatly aids in this regularization.

Huq et al. [89, 90] choose, instead, to use a completely different computation technique for Θ , based on *3-dimensional Cartesian* finite differencing:

1. They observe that the scalar field F defined by Equation (7) can be evaluated at any (3-dimensional) position in the slice by computing the corresponding (r, θ, ϕ) using the usual flat-space formulas, then interpolating h in the 2-dimensional (θ, ϕ) surface grid.
2. Rewrite the apparent horizon condition (15) in terms of F and its (3-dimensional) *Cartesian* derivatives,

$$\Theta \equiv \Theta(F, \partial_i F, \partial_{ij} F; g_{ij}, \partial_k g_{ij}, K_{ij}) = 0. \quad (27)$$

Huq et al. [89, 90] give the $\Theta(F, \partial_i F, \partial_{ij} F)$ function explicitly.

3. For each (latitude/longitude) grid point on the trial horizon surface, define a $3 \times 3 \times 3$ -point *local Cartesian grid* centered at that point. The spacing of this grid should be such as to allow accurate finite differencing, i.e. in practice it should probably be roughly comparable to that of the underlying numerical-relativity simulation’s grid.
4. Evaluate F on the local Cartesian grid as described in Step 1 above.

³⁹See [133] for a substantially revised version of [132].

5. Evaluate the Cartesian derivatives in Equation (27) by centered 2nd order Cartesian finite differencing of the F values on the local Cartesian grid.

Comparing the different ways of evaluating Θ , 2-dimensional angular finite differencing of Equation (16) seems to me to be both simpler (easier to program) and likely more efficient than 3-dimensional Cartesian finite differencing of Equation (27).

8.5.3 Solving the nonlinear elliptic PDE

A variety of algorithms are possible for actually solving the nonlinear elliptic PDE (16) (or (27) for the Huq et al. [89, 90] horizon finder).

The most common choice is to use some variant of Newton’s method. That is, finite differencing Equation (16) or (27) (as appropriate) gives a system of N_{ang} nonlinear algebraic equations for the horizon shape function h at the N_{ang} angular grid points; these can be solved by Newton’s method in N_{ang} dimensions. (As explained by Thornburg [153, Section VIII.C], this is usually equivalent to applying the Newton–Kantorovich algorithm [37, Appendix C] to the original nonlinear elliptic PDE (16) or (27).)

Newton’s method converges very quickly once the trial horizon surface is sufficiently close to a solution (a MOTS). However, for a less accurate initial guess, Newton’s method may converge very slowly or even fail to converge at all. There is no usable way of determining *a priori* just how large the radius of convergence of the iteration will be, but in practice $\frac{1}{4}$ to $\frac{1}{3}$ of the horizon radius is often a reasonable estimate⁴⁰.

Thornburg [153] described the use of various “line search” modifications to Newton’s method to improve its radius and robustness of convergence, and reported that even fairly simple modifications of this sort roughly doubled the radius of convergence.

Schnetter [132, 133] used the PETSC general-purpose elliptic-solver library [22, 23, 24] to solve the equations. This offers a wide variety of Newton-like algorithms already implemented in a highly optimized form.

Rather than Newton’s method or one of its variants, Shibata et al. [146, 147] use a functional-iteration algorithm directly on the nonlinear elliptic PDE (16). This seems likely to be less efficient than Newton’s method but avoids having to compute and manipulate the Jacobian matrix.

8.5.4 The Jacobian matrix

Newton’s method, and all its variants, require an explicit computation of the Jacobian matrix

$$\mathbf{J}_{\text{IJ}} = \frac{\partial \Theta_{\text{I}}}{\partial h_{\text{J}}}, \quad (28)$$

where the indices I and J label angular grid points on the horizon surface (or equivalently on S^2).

Notice that \mathbf{J} includes contributions both from the direct dependence of Θ on h , $\partial_u h$, and $\partial_{uv} h$, and also from the indirect dependence of Θ on h through the position-dependence of the geometry variables g_{ij} , $\partial_k g_{ij}$, and K_{ij} (since Θ depends on the geometry variables *at the horizon surface position*, and this position is determined by h). Thornburg [153] discusses this indirect dependence in detail.

There are two basic ways to compute the Jacobian matrix.

⁴⁰Thornburg [153] used a Monte-Carlo survey of horizon-shape perturbations to quantify the radius of convergence of Newton’s method for apparent horizon finding. He found that if strong high-spatial-frequency perturbations are present in the slice’s geometry then the radius of convergence may be very small. Fortunately, this problem rarely occurs in practice.

Numerical Perturbation:

The simplest way to determine the Jacobian matrix is by “numerical perturbation”, where for each horizon-surface grid point j , h is perturbed by some (small) amount ε at the j th grid point (that is, $h_I \rightarrow h_I + \varepsilon\delta_{IJ}$), and the expansion Θ is recomputed⁴¹. The j th column of the Jacobian matrix (28) is then estimated as

$$\mathbf{J}_{IJ} \approx \frac{\Theta_I(h + \varepsilon\delta_{IJ}) - \Theta_I(h)}{\varepsilon}. \quad (29)$$

Curtis and Reid [53], and Stoer and Bulirsch [150, Section 5.4.3] discuss the optimum choice of ε in this algorithm⁴².

This algorithm is easy to program but somewhat inefficient. It is used by a number of researchers including Schnetter [132, 133], and Huq et al. [89, 90].

Symbolic Differentiation:

A more efficient, although somewhat more complicated, way to determine the Jacobian matrix is the “symbolic differentiation” algorithm described by Thornburg [153] and also used by Pasch [118], Shibata et al. [146, 147], and Thornburg [156]. Here the internal structure of the finite differenced $\Theta(h)$ function is used to directly determine the Jacobian matrix elements.

This algorithm is best illustrated by an example which is simpler than the full apparent horizon equation: Consider the flat-space Laplacian in standard (θ, ϕ) polar-spherical coordinates,

$$\Delta h \equiv \partial_{\theta\theta} h + \frac{\partial_{\theta} h}{\tan \theta} + \frac{\partial_{\phi\phi} h}{\sin^2 \theta}. \quad (30)$$

Suppose we discretize this with centered 2nd order finite differences in θ and ϕ . Then neglecting finite-differencing truncation errors, and temporarily adopting the usual notation for 2-dimensional grid functions, $h_{i,j} \equiv h(\theta=\theta_i, \phi=\phi_j)$, our discrete approximation to Δh is given by

$$(\Delta h)_{i,j} = \frac{h_{i-1,j} - 2h_{i,j} + h_{i+1,j}}{(\Delta\theta)^2} + \frac{1}{\tan \theta} \frac{h_{i+1,j} - h_{i-1,j}}{2\Delta\theta} + \frac{1}{\sin^2 \theta} \frac{h_{i,j-1} - 2h_{i,j} + h_{i,j+1}}{(\Delta\phi)^2}. \quad (31)$$

The Jacobian of Δh is thus given by

$$\frac{\partial(\Delta h)_{(i,j)}}{\partial h_{(k,\ell)}} = \begin{cases} \frac{1}{(\Delta\theta)^2} \pm \frac{1}{2 \tan \theta \Delta\theta} & \text{if } (k, \ell) = (i \pm 1, j), \\ \frac{1}{\sin^2 \theta (\Delta\phi)^2} & \text{if } (k, \ell) = (i, j \pm 1), \\ -\frac{2}{(\Delta\theta)^2} - \frac{2}{\sin^2 \theta (\Delta\phi)^2} & \text{if } (k, \ell) = (i, j), \\ 0 & \text{otherwise.} \end{cases} \quad (32)$$

Thornburg [153] describes how to generalize this to nonlinear differential operators without having to explicitly manipulate the nonlinear finite difference equations.

⁴¹A very important optimization here is that Θ only needs to be recomputed within the finite-difference domain of dependence of the j th grid point.

⁴²Because of the one-sided finite differencing, the approximation (29) is only $\mathcal{O}(\varepsilon)$ accurate. However, in practice this does not seriously impair the convergence of a horizon finder, and the extra cost of a centered-finite-differencing $\mathcal{O}(\varepsilon^2)$ approximation is not warranted.

8.5.5 Solving the linear equations

All the algorithms described in Section 8.5.3 for treating nonlinear elliptic PDEs require solving a sequence of linear systems of N_{ang} equations in N_{ang} unknowns. N_{ang} is typically on the order of a few thousand, and the Jacobian matrices in question are sparse due to the locality of the angular finite differencing (see Section 8.5.4). Thus, for reasonable efficiency, it is essential to use linear solvers that exploit this sparsity. Unfortunately, many such algorithms/codes only handle symmetric positive-definite matrices while, due to the angular boundary conditions⁴³ (see Section 8.5.1), the Jacobian matrices that arise in apparent horizon finding are generally neither of these.

The numerical solution of large sparse linear systems is a whole subfield of numerical analysis. See, for example, Duff, Erisman, and Reid [65], and Saad [130] for extensive discussions⁴⁴. In practice, a numerical relativist is unlikely to write her own linear solver but, rather, will use an existing subroutine (library).

Kershaw's [94] ILUCG iterative solver is often used; this is only moderately efficient, but is quite easy to program⁴⁵. Schnetter [132, 133] reports good results with an ILU-preconditioned GMRES solver from the PETSC library. Thornburg [156] experimented with both an ILUCG solver and a direct sparse LU decomposition solver (Davis' UMFPACK library [57, 58, 56, 55, 54]), and found each to be more efficient in some situations; overall, he found the UMFPACK solver to be the best choice.

8.5.6 Sample results

As an example of the results obtained with this type of apparent horizon finder, Figure 10 shows the numerically-computed apparent horizons (actually, MOTSs) at two times in a head-on binary black hole collision. (The physical system being simulated here is very similar to that simulated by Matzner et al. [108], a view of whose event horizon is shown in Figure 5.)

As another example, Figure 11 shows the time dependence of the irreducible masses of apparent horizons found in a (spiraling) binary black hole collision, simulated at several different grid resolutions, as found by both AHFINDERDIRECT and another CACTUS apparent horizon finder, AHFINDER⁴⁶. For this evolution, the two apparent horizon finders give irreducible masses which agree to within about 2% for the individual horizons and 0.5% for the common horizon.

As a final example, Figure 8 shows the numerically-computed event and apparent horizons in the collapse of a rapidly rotating neutron star to a Kerr black hole. (The event horizons were computed using the EHFINDER code described in Section 5.3.3.)

8.5.7 Summary of elliptic-PDE algorithms/codes

Elliptic-PDE apparent horizon finders have been developed by many researchers, including Eardley [67]⁴⁷, Cook [50, 52, 51], and Thornburg [153] in axisymmetry, and Shibata et al. [146, 147], Huq et al. [89, 90], Schnetter [132, 133], and Thornburg [156] in fully generic slices.

⁴³Or the interpatch interpolation conditions in Thornburg's multiple-grid-patch scheme [156].

⁴⁴Multigrid algorithms are also important here; these exploit the geometric structure of the underlying elliptic PDE. See Briggs, Henson, and McCormick [42] and Trottenberg, Oosterlee, and Schüller [158] for general introductions to multigrid algorithms.

⁴⁵Madderom's Fortran subroutine DILUCG [107], which implements the method of [94], has been used by a number of numerical relativists for both this and other purposes.

⁴⁶AHFINDER incorporates both a minimization algorithm (Section 8.3) and a fast-flow algorithm (Section 8.7.4); these tests used the fast-flow algorithm.

⁴⁷This paper does not say how the author finds apparent horizons, but [68, page 135] cites a preprint of this as treating the apparent-horizon equation as a 2-point (ODE) boundary value problem: Eardley uses a 'beads on a string' technique to solve the set of simultaneous equations, i.e., imagining the curve to be defined as a bead on each ray of constant angle. He solves for the positions on each ray at which the relation is satisfied everywhere.




Figure 10: *This figure shows the numerically-computed apparent horizons (actually MOTSs) at two times in a head-on binary black hole collision. The black holes are colliding along the z axis. Figure reprinted with permission from [156]. © 2004 by IOP Publishing Ltd.*




Figure 11: *This figure shows the irreducible masses ($\sqrt{\text{area}}/(16\pi)$) of individual and common apparent horizons in a binary black hole collision, as calculated by two different apparent horizon finders in the CACTUS toolkit, AHFINDER and AHFINDERDIRECT. (AHFINDERDIRECT was also run in simulations at several different resolutions.) Notice that when both apparent horizon finders are run in the same simulation (resolution $dx = 0.080$), there are only small differences between their results. Figure reprinted with permission from [5]. © 2005 by the American Physical Society.*

Elliptic-PDE algorithms are (or can be implemented to be) among the fastest horizon finding algorithms. For example, running on a 1.7 GHz processor, Thornburg’s AHFINDERDIRECT [156] averaged 1.7 s per horizon finding, as compared with 61 s for an alternate “fast-flow” apparent horizon finder AHFINDER (discussed in more detail in Section 8.7)⁴⁸. However, achieving maximum performance comes at some cost in implementation effort (e.g. the “symbolic differentiation” Jacobian computation discussed in Section 8.5.4).

Elliptic-PDE algorithms are probably somewhat more robust in their convergence (i.e. they have a slightly larger radius of convergence) than other types of local algorithms, particularly if the “line search” modifications of Newton’s method described by Thornburg [153] are implemented⁴⁹. Their typical radius of convergence is on the order of $\frac{1}{3}$ of the horizon radius, but cases are known where it is much smaller. For example, Schnetter, Herrmann, and Pollney [135] report that (with no “line search” modifications) it is only about 10% for some slices in a binary black hole coalescence simulation.

Schnetter’s TGRAPPARENTHORIZON2D [132, 133] and Thornburg’s AHFINDERDIRECT [156] are both elliptic-PDE apparent horizon finders implemented as freely available modules (“thorns”) in the CACTUS computational toolkit (see Table 2). Both work with either the PUGH unigrid driver or the CARPET mesh-refinement driver for CACTUS. TGRAPPARENTHORIZON2D is no longer maintained, but AHFINDERDIRECT is actively supported and is now used by many different research groups⁵⁰.

8.6 Horizon pretracking

Schnetter et al. [133, 135] introduced the important concept of “horizon pretracking”. They focus on the case where we want to find a common apparent horizon as soon as it appears in a binary black-hole (or neutron-star) simulation. While a global (flow) algorithm (Section 8.7) could be used to find this common apparent horizon, these algorithms tend to be very slow. They observe that the use of a local (elliptic-PDE) algorithm for this purpose is somewhat problematic:

The common [apparent] horizon [...] appears instantaneously at some late time and without a previous good guess for its location. In practice, an estimate of the surface location and shape can be put in by hand. The quality of this guess will determine the rate of convergence of the finder and, more seriously, also determines whether a horizon is found at all. Gauge effects in the strong field region can induce distortions that have a large influence on the shape of the common horizon, making them difficult to predict, particularly after a long evolution using dynamical coordinate conditions. As such, it can be a matter of some expensive trial and error to find the common apparent horizon at the earliest possible time. Further, if a common apparent horizon is not found, it is not clear whether this is because there is none, or whether there exists one which has only been missed due to unsuitable initial guesses – for a fast apparent horizon finder, a good initial guess is crucial.

Pretracking tries (usually successfully) to eliminate these difficulties by determining – *before* it appears – approximately where (in space) and when (in time) the common apparent horizon will appear.

⁴⁸As another comparison, the LORENE apparent horizon finder (discussed in more detail in Section 8.4.2), running on a roughly similar processor, takes between 3 and 6 seconds to find apparent horizons to comparable accuracy.

⁴⁹The convergence problems, which Thornburg [153] noted when high-spatial-frequency perturbations are present in the slice’s geometry, seem to be rare in practice.

⁵⁰In addition, at least two different research groups have now ported, or are in the process of porting, AHFINDERDIRECT to their own (non-CACTUS) numerical relativity codes.

8.6.1 Constant-expansion surfaces

The basic idea of horizon pretracking is to consider surfaces of constant expansion (“CE surfaces”), i.e. smooth closed orientable 2-surfaces in a slice satisfying the condition

$$\Theta = E, \quad (33)$$

where the expansion E is a specified real number. Each marginally outer trapped surface (including the apparent horizon) is thus a CE surface with expansion $E = 0$; more generally Equation(33) defines a 1-parameter family of 2-surfaces in the slice. As discussed by Schnetter et al. [133, 135], for asymptotically flat slices containing a compact strong-field region, some of the $E > 0$ members of this family typically foliate the weak-field region.

In the binary-coalescence context, for each $t = \text{constant}$ slice we define E_* to be the smallest $E \geq 0$ for which a CE surface (containing both strong-field regions) exists with expansion E . If $E_* = 0$ this “minimum-expansion CE surface” is the common apparent horizon, while if $E_* > 0$ this surface is an approximation to where the common apparent horizon *will* appear. We expect the minimum-expansion CE surface to change continuously during the evolution and its expansion E_* to decrease towards 0. Essentially, horizon pretracking follows the time evolution of the minimum-expansion CE surface and uses it as an initial guess for (searching for) the common apparent horizon.

8.6.2 Generalized constant-expansion surfaces

Schnetter [133] implemented an early form of horizon pretracking, which followed the evolution of the minimum-expansion constant-expansion surface, and found that it worked well for simple test problems. However, Schnetter et al. [135] found that for more realistic binary-black-hole coalescence systems the algorithm needs to be extended:




Figure 12: This figure shows the expansion Θ (left scale), and the “generalized expansions” $r\Theta$ (left scale) and $r^2\Theta$ (right scale), for various $r = \text{constant}$ surfaces in an Eddington–Finkelstein slice of Schwarzschild spacetime. Notice that all three functions have zeros at the horizon $r = 2m$, and that while Θ has a maximum at $r \approx 4.4m$, both $r\Theta$ and $r^2\Theta$ increase monotonically with r .

- While the expansion is zero for a common apparent horizon, it is also zero for a 2-sphere at spatial infinity. Figure 12 illustrates this for Schwarzschild spacetime. Notice that for small

positive E_* there will generally be two distinct CE surfaces with $E = E_*$, an inner surface just outside the horizon and an outer one far out in the weak-field region. The inner CE surface converges to the common apparent horizon as E_* decreases towards 0; this surface is the one we would like the pretracking algorithm to follow. Unfortunately, without measures such as those described below, there is nothing to prevent the algorithm from following the outer surface, which does *not* converge to the common apparent horizon as E_* decreases towards 0.

- In a realistic binary-coalescence simulation, the actual minimum-expansion CE surface may be highly distorted and, thus, hard to represent accurately with a finite-resolution angular grid.

Schnetter et al. [135] discuss these problems in more detail, arguing that to solve them, the expansion Θ should be generalized to a “shape function” H given by one of

$$\begin{aligned} H_1 &= \Theta, \\ H_r &= h\Theta, \\ H_{r,2} &= h^2\Theta, \end{aligned} \tag{34}$$

CE surfaces are then generalized to surfaces satisfying

$$H = E \tag{35}$$

for some specified $E \geq 0$.

Note that unlike H_1 , both H_r and $H_{r,2}$ are typically monotonic with radius. Neither H_r nor $H_{r,2}$ are 3-covariantly defined, but they both still have the property that $E = 0$ in Equation (35) implies the surface is a MOTS, and in practice they work better for horizon pretracking.

8.6.3 Goal functions

To define the single “smallest” surface at each time, Schnetter et al. [135] introduce a second generalization, that of a “goal function” G , which maps surfaces to real numbers. The pretracking search then attempts, on each time slice, to find the surface (shape) satisfying $H = E$ with the minimum value of G . They experimented with several different goal functions,

$$\begin{aligned} G_H &= \overline{H}, \\ G_{rH} &= \overline{hH}, \\ G_r &= \overline{h}, \end{aligned} \tag{36}$$

where in each case the overbar ($\overline{\quad}$) denotes an average over the surface⁵¹.

8.6.4 The pretracking search

Schnetter’s [133] original implementation of horizon pretracking (which followed the evolution of the minimum-expansion CE surface) used a binary search on the desired expansion E . Because E appears only on the right hand side of the generalized CE condition (35), it is trivial to modify any apparent horizon finder to search for a surface of specified expansion E . (Schnetter used his

⁵¹Schnetter et al. [135] use a simple arithmetic mean over all surface grid points. In theory this average could be defined 3-covariantly by taking the induced metric on the surface into account, but in practice they found that this was not worth the added complexity.

TGRAPPARENTHORIZON2D elliptic-PDE apparent horizon finder described in Section 8.5.7 for this.) A binary search on E can then be used to find the minimum value E_* .⁵²

Implementing a horizon pretracking search on any of the generalized goal functions (36) is conceptually similar but somewhat more involved: As described by Schnetter et al. [135] for the case of an elliptic-PDE apparent horizon finder⁵³, we first write the equation defining a desired pretracking surface as

$$H - \bar{H} + G - p = 0, \quad (37)$$

where p is the desired value of the goal function G . Since H is the only term in Equation (37) which varies over the surface, it must be constant for the equation to be satisfied. In this case $H - \bar{H}$ vanishes, so the equation just gives $G = p$, as desired.

Because \bar{H} depends on H at *all* surface points, directly finite differencing Equation (37) would give a non-sparse Jacobian matrix, which would greatly slow the linear-solver phase of the elliptic-PDE apparent horizon finder (Section 8.5.5). Schnetter et al. [135, Section III.B] show how this problem can be solved by introducing a single extra unknown into the discrete system. This gives a Jacobian which has a single non-sparse row and column, but is otherwise sparse, so the linear equations can still be solved efficiently.

When doing the pretracking search, the cost of a single binary-search iteration is approximately the same as that of finding an apparent horizon. Schnetter et al. [135, Figure 5] report that their pretracking implementation (a modified version of Thornburg’s AHFINDERDIRECT [156] elliptic-PDE apparent horizon finder described in Section 8.5.7) typically takes on the order of 5 to 10 binary-search iterations⁵⁴. The cost of pretracking is thus on the order of 5 to 10 times that of finding a single apparent horizon. This is substantial, but not prohibitive, particularly if the pretracking algorithm is not run at every time step.

8.6.5 Sample results

As an example of the results obtained from horizon pretracking, Figure 13 shows the expansion Θ for various pretracking surfaces (i.e. various choices for the shape function H in a head-on binary black hole collision). Notice how all three of the shape functions (34) result in pretracking surfaces whose expansions converge smoothly to zero just when the apparent horizon appears (at about $t = 1.1$).

As a further example, Figure 14 shows the pretracking surfaces (more precisely, their cross sections projected into the black holes’ orbital plane) at various times in a spiraling binary black hole collision.

8.6.6 Summary of horizon pretracking

Pretracking is a very valuable addition to the horizon finding repertoire: It essentially gives a local algorithm (in this case, an elliptic-PDE algorithm) most of the robustness of a global algorithm in

⁵²There is one complication here: Any local apparent horizon finding algorithm may fail if the initial guess is not good enough, *even if the desired surface is actually present*. The solution is to use the constant-expansion surface for a slightly larger expansion E as an initial guess, gradually “walking down” the value of E to find the minimum value E_* . Thornburg [156, Appendix C] describes such a “continuation-algorithm binary search” algorithm in detail.

⁵³As far as I know this is the only case that has been considered for horizon pretracking. Extension to other types of apparent horizon finders might be a fruitful area for further research.

⁵⁴This refers to the period before a common apparent horizon is found. Once a common apparent horizon is found, then pretracking can be disabled as the apparent horizon finder can easily “track” the apparent horizon’s motion from one time step to the next. With a binary search the number of iterations depends only weakly (logarithmically) on the pretracking algorithm’s accuracy tolerance. It might be possible to replace the binary search by a more sophisticated 1-dimensional search algorithm (I discuss such algorithms in Appendix A), potentially cutting the number of iterations substantially. This might be a fruitful area for further research.




Figure 13: This figure shows the expansion Θ for various pretracking surfaces, i.e. for various choices for the shape function H , in a head-on binary black hole collision. Notice how the three shape functions (34) (here labelled Θ , $r\Theta$, and $r^2\Theta$) result in pretracking surfaces whose expansions converge smoothly to zero just when the apparent horizon appears (at about $t = 1.1$). Notice also that these three expansions have all converged to each other somewhat before the common apparent horizon appears. Figure reprinted with permission from [135]. © 2005 by the American Physical Society.

terms of finding a common apparent horizon as soon as it appears. It is implemented as a higher-level algorithm which uses a slightly-modified elliptic-PDE apparent horizon finding algorithm as a “subroutine”.

The one significant disadvantage of pretracking is its cost: Each pretracking search typically takes 5 to 10 times as long as finding an apparent horizon. Further research to reduce the cost of pretracking would be useful.

Schnetter et al.’s pretracking implementation [135] is implemented as a set of modifications to Thornburg’s AHFINDERDIRECT [156] apparent horizon finder. Like the original AHFINDERDIRECT, the modified version is a freely available “thorn” in the CACTUS toolkit (see Table 2).

8.7 Flow algorithms

Flow algorithms define a “flow” on 2-surfaces, i.e. they define an evolution of 2-surfaces in some pseudo-time λ , such that the apparent horizon is the $\lambda \rightarrow \infty$ limit of a (any) suitable starting surface. Flow algorithms are different from other apparent horizon finding algorithms (except for zero-finding in spherical symmetry) in that their convergence does not depend on having a good initial guess. In other words, flow algorithms are global algorithms (Section 7.7).

To find the (an) apparent horizon, i.e. an *outermost* MOTS, the starting surface should be outside the largest possible MOTS in the slice. In practice, it generally suffices to start with a 2-sphere of areal radius substantially greater than $2m_{\text{ADM}}$.

The global convergence property requires that a flow algorithm always flow from a large starting surface into the apparent horizon. This means that the algorithm gains no particular benefit from already knowing the approximate position of the apparent horizon. In particular, flow algorithms are no faster when “tracking” the apparent horizon (repeatedly finding it at frequent intervals) in a numerical time evolution. (In contrast, in this situation a local apparent horizon finding algorithm

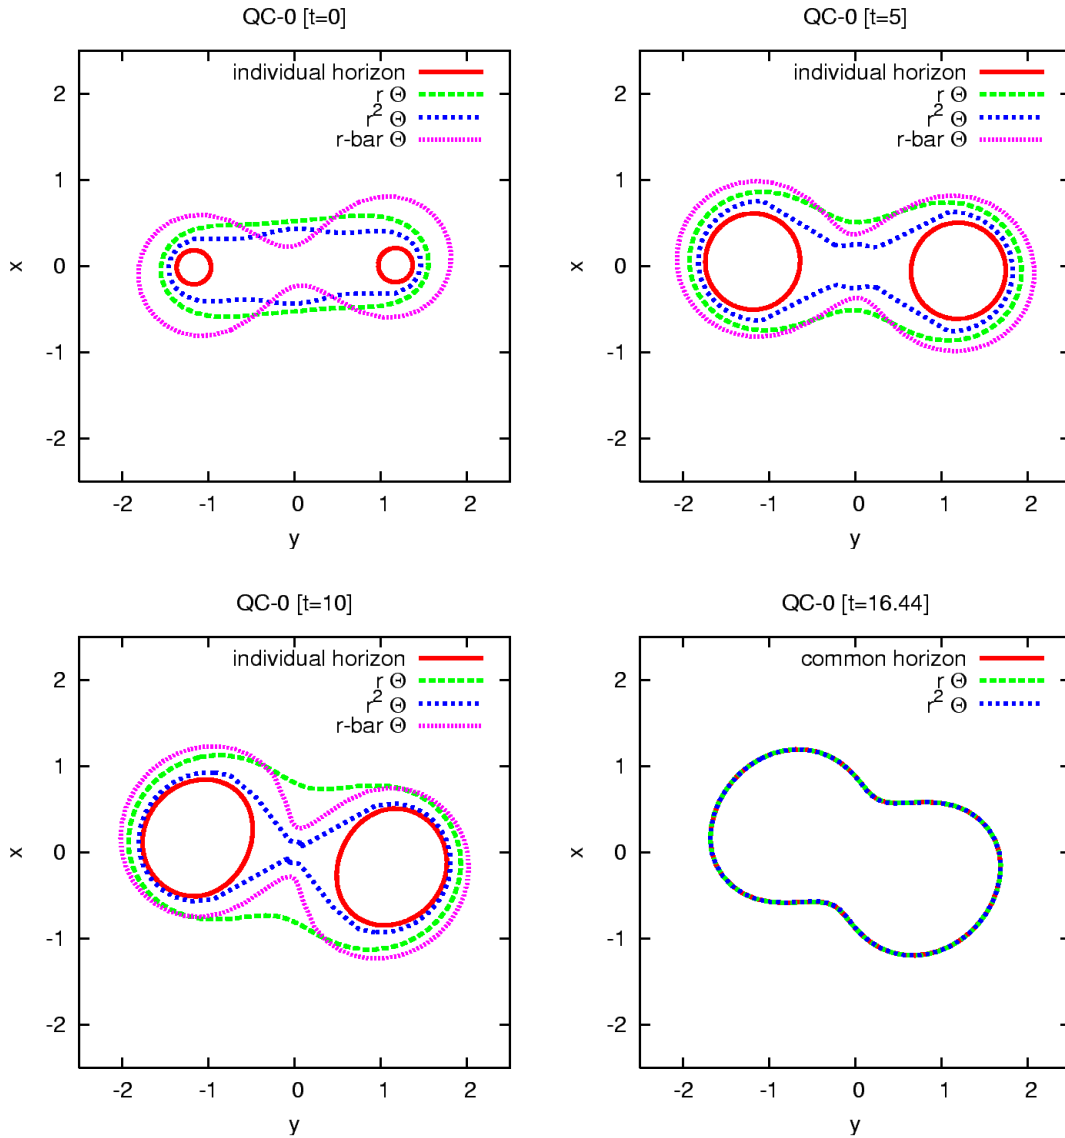


Figure 14: *This figure shows the pretracking surfaces at various times in a spiraling binary black hole collision, projected into the black holes' orbital plane. (The apparent slow drift of the black holes in a clockwise direction is an artifact of the corotating coordinate system; the black holes are actually orbiting much faster, in a counterclockwise direction.) Notice how, even well before the common apparent horizon first appears ($t = 16.44 m_{\text{ADM}}$, bottom right plot), the r_{Θ} pretracking surface is already a reasonable approximation to the eventual common apparent horizon's shape. Figure reprinted with permission from [135]. © 2005 by the American Physical Society.*

can use the most recent previously-found apparent horizon as an initial guess, greatly speeding the algorithm's convergence⁵⁵).

Flow algorithms were first proposed for apparent horizon finding by Tod [157]. He initially considered the case of a time-symmetric slice (one where $K_{ij} = 0$). In this case, a MOTS (and thus an apparent horizon) is a surface of minimal area and may be found by a “mean curvature flow”

$$\partial_\lambda x^i = -\kappa s^i, \quad (38)$$

where x^i are the spatial coordinates of a horizon-surface point, s^i is as before the outward-pointing unit 3-vector normal to the surface, and $\kappa \equiv \nabla_k s^k$ is the mean curvature of the surface as embedded in the slice. This is a gradient flow for the surface area, and Grayson [79] has proven that if the slice contains a minimum-area surface, this will in fact be the stable $\lambda \rightarrow \infty$ limit of this flow. Unfortunately, this proof is valid only for the time-symmetric case.

For non-time-symmetric slices, Tod [157] proposed generalizing the mean curvature flow to the “expansion flow”

$$\partial_\lambda x^i = -\Theta s^i. \quad (39)$$

There is no theoretical proof that this flow will converge to the (an) apparent horizon, but several lines of argument make this convergence plausible:

- The expansion flow is identical to the mean curvature flow (38) in the principal part.
- The expansion flow's velocity is clearly zero on an apparent horizon.
- More generally, a simple argument due to Bartnik [25]⁵⁶ shows that the expansion flow can never move a (smooth) test surface *through* an apparent horizon. Suppose, to the contrary, that the test surface \mathcal{T} is about to move through an apparent horizon \mathcal{H} , i.e. since both surfaces are by assumption smooth, that \mathcal{T} and \mathcal{H} touch at single (isolated) point P . At that point, \mathcal{T} and \mathcal{H} obviously have the same g_{ij} and K_{ij} , and they also have the same s^i (because P is isolated). Hence the only term in Θ (as defined by Equation (15)) which differs between \mathcal{T} and \mathcal{H} is $\nabla_i s^i$. Clearly, if \mathcal{T} is outside \mathcal{H} and they touch at the single isolated point P , then relative to \mathcal{H} , \mathcal{T} must be concave outwards at P , so that $\nabla_i s^i(\mathcal{T}) < \nabla_i s^i(\mathcal{H})$. Thus the expansion flow (39) will move \mathcal{T} outwards, *away* from the apparent horizon. (If \mathcal{T} lies inside \mathcal{H} the same argument holds with signs reversed appropriately.)

Numerical experiments by Bernstein [28], Shoemaker et al. [148, 149], and Pasch [118] show that in practice the expansion flow (39) does in fact converge robustly to the apparent horizon.

In the following I discuss a number of important implementation details for, and refinements of, this basic algorithm.

8.7.1 Implicit pseudo-time stepping

Assuming the Strahlkörper surface parameterization (4), the expansion flow (39) is a *parabolic* equation for the horizon shape function h .⁵⁷ This means that any fully explicit scheme to integrate it (in the pseudo-time λ) must severely restrict its pseudo-time step $\Delta\lambda$ for stability, and this restriction grows (quadratically) worse at higher spatial resolutions⁵⁸. This makes the horizon finding process very slow.

⁵⁵Alternatively, a flow algorithm could use the most recent previously-found apparent horizon as an initial guess. In this case the algorithm would have only local convergence (in particular, it would probably fail to find a new outermost MOTS that appeared well outside the previously-found MOTS). However, the algorithm would only need to flow the surface a small distance, so the algorithm should be fairly fast.

⁵⁶Cited as Ref. [17] by [80].

⁵⁷Linearizing the $\Theta(h)$ function (16) gives a negative Laplacian in h as the principal part.

⁵⁸For a spatial resolution Δx , an explicit scheme is generally limited to a pseudo-time step $\Delta\lambda \lesssim (\Delta x)^2$.

To avoid this restriction, practical implementations of flow algorithms use implicit pseudo-time integration schemes; these can have large pseudo-time steps and still be stable. Because we only care about the $\lambda \rightarrow \infty$ limit, a highly accurate pseudo-time integration is not important; only the accuracy of approximating the spatial derivatives matters. Bernstein [28] used a modified Du Fort–Frankel scheme [64]⁵⁹ but found some problems with the surface shape gradually developing high-spatial-frequency noise. Pasch [118] reports that an “exponential” integrator (Hochbruck et al. [85]) works well, provided the flow’s Jacobian matrix is computed accurately⁶⁰. The most common choice is probably that of Shoemaker et al. [148, 149], who use the iterated Crank–Nicholson (“ICN”) scheme⁶¹. They report that this works very well; in particular, they do not report any noise problems.

By refining his finite-element grid (Section 2.3) in a hierarchical manner, Metzger [109] is able to use standard conjugate-gradient elliptic solvers in a multigrid-like fashion⁶², using each refinement level’s solution as an initial guess for the next higher refinement level’s iterative solution. This greatly speeds the flow integration: Metzger reports that the performance of the overall surface-finding algorithm is “of the same order of magnitude” as that of Thornburg’s AHFIND-ERDIRECT [156] elliptic-PDE apparent horizon finder (described in Section 8.5.7).

In a more general context than numerical relativity, Osher and Sethian [116] have discussed a general class of numerical algorithms for integrating “fronts propagating with curvature-dependent speed”. These flow a level-set function (Section 2.1) which implicitly locates the actual “front”.

8.7.2 Varying the flow speed

Another important performance optimization of the standard expansion flow (39) is to replace Θ in the right-hand side by a suitable nonlinear function of Θ , chosen so the surface shrinks faster when it is far from the apparent horizon. For example, Shoemaker et al. [148, 149] use the flow

$$\partial_\lambda x^i = - \left[(\Theta - c) \arctan^2 \left(\frac{\Theta - c}{\Theta_0} \right) \right] s^i \quad (40)$$

for this purpose, where Θ_0 is the value of Θ on the initial-guess surface, and c (which is gradually decreased towards 0 as the iteration proceeds) is a “goal” value for Θ .

8.7.3 Surface representation and the handling of bifurcations

Since a flow algorithm starts with (topologically) a single large 2-sphere, if there are multiple apparent horizons present the surface must change topology (bifurcate) at some point in the flow. Depending on how the surface is represented, this may be easy or difficult.

Pasch [118] and Shoemaker et al. [148, 149] use a level-set function approach (Section 2.1). This automatically handles any topology or topology change. However, it has the drawback of requiring the flow to be integrated throughout the entire volume of the slice (or at least in some neighborhood of each surface). This is likely to be much more expensive than only integrating the flow on the surface itself. Shoemaker et al. also generate an explicit Strahlkörper surface representation (Section 2.2), monitoring the surface shape to detect an imminent bifurcation and reparameterizing the shape into 2 separate surfaces if a bifurcation happens.

⁵⁹Richtmyer and Morton [129, Section 7.5] give a very clear presentation and analysis of the Du Fort–Frankel scheme.

⁶⁰More precisely, Pasch [118] found that that an exponential integrator worked well when the flow’s Jacobian matrix was computed exactly (using the symbolic-differentiation technique described in Section 8.5.4). However, when the Jacobian matrix was approximated using the numerical-perturbation technique described in Section 8.5.4, Pasch found that the pseudo-time integration became unstable at high numerical resolutions. Pasch [118] also notes that the exponential integrator uses a very large amount of memory.

⁶¹Teukolsky [152], and Leiler and Rezzolla [101] have analyzed ICN’s stability under various conditions.

⁶²See [42, 158] for general introductions to multigrid algorithms for elliptic PDEs.

Metzger [109] uses a finite-element surface representation (Section 2.3), which can represent any topology. However, if the flow bifurcates, then to explicitly represent each apparent horizon the code must detect that the surface self-intersects, which may be expensive.

8.7.4 Gundlach’s “fast flow”

Gundlach [80] introduced the important concept of a “fast flow”. He observed that the subtraction and inversion of the flat-space Laplacian in the Nakamura–Kojima–Oohara spectral integral-iteration algorithm (Section 8.4) is an example of “a standard way of solving nonlinear elliptic problems numerically, namely subtracting a simple linear elliptic operator from the nonlinear one, inverting it by pseudo-spectral algorithms and iterating”. Gundlach then interpreted the Nakamura–Kojima–Oohara algorithm as a type of flow algorithm where each pseudo-time step of the flow corresponds to a single functional-iteration step of the Nakamura–Kojima–Oohara algorithm.

In this framework, Gundlach defines a 2-parameter family of flows interpolating between the Nakamura–Kojima–Oohara algorithm and Tod’s [157] expansion flow (39),

$$\partial_\lambda h = -A(1 - B\Delta)^{-1}\rho\Theta, \quad (41)$$

where $A \geq 0$ and $B \geq 0$ are parameters, $\rho > 0$ is a weight functional which depends on h through at most 1st derivatives, Δ is the flat-space Laplacian operator, and $(1 - B\Delta)^{-1}$ denotes inverting the operator $(1 - B\Delta)$. Gundlach then argues that intermediate “fast flow” members of this family should be useful compromises between the speed of the Nakamura–Kojima–Oohara algorithm and the robustness of Tod’s expansion flow. Based on numerical experiments, Gundlach suggests a particular choice for the weight functional ρ and the parameters A and B . The resulting algorithm updates high-spatial-frequency components of h essentially the same as the Nakamura–Kojima–Oohara algorithm but should reduce low-spatial-frequency error components faster. Gundlach’s algorithm also completely avoids the need for numerically solving Equation (23) for the a_{00} coefficient in the Nakamura–Kojima–Oohara algorithm.

Alcubierre’s AHFINDER [4] horizon finder includes an implementation of Gundlach’s fast flow algorithm⁶³. AHFINDER is implemented as a freely available module (“thorn”) in the CACTUS computational toolkit (see Table 2) and has been used by many research groups.

8.7.5 Summary of flow algorithms/codes

Flow algorithms are the only truly global apparent horizon finding algorithms and, as such, can be much more robust than local algorithms. In particular, flow algorithms can guarantee convergence to the *outermost* MOTS in a slice. Unfortunately, these convergence guarantees hold only for time-symmetric slices.

In the forms which have strong convergence guarantees, flow algorithms tend to be very slow. (Metzger’s algorithm [109] is a notable exception: It is very fast.) There are modifications which can make flow algorithms much faster, but then their convergence is no longer guaranteed. In particular, practical experience has shown that in some binary black hole coalescence simulations (Alcubierre et al. [5], Diener et al. [62]), “fast flow” algorithms (Section 8.7.4) can miss common apparent horizons which are found by other (local) algorithms.

Alcubierre’s apparent horizon finder AHFINDER [4] includes a “fast flow” algorithm based on the work of Gundlach [80]⁶³. It is implemented as a freely available module (“thorn”) in the CACTUS computational toolkit (see Table 2) and has been used by a number of research groups.

⁶³AHFINDER also includes a minimization algorithm (Section 8.3).

9 Summary of Algorithms/Codes for Finding Apparent Horizons

There are many apparent horizon finding algorithms, with differing trade-offs between speed, robustness of convergence, accuracy, and ease of programming.

In spherical symmetry, zero-finding (Section 8.1) is fast, robust, and easy to program. In axisymmetry, shooting algorithms (Section 8.2) work well and are fairly easy to program. Alternatively, any of the algorithms for generic slices (summarized below) can be used with implementations tailored to the axisymmetry.

Minimization algorithms (Section 8.3) are fairly easy to program, but when the underlying simulation uses finite differencing these algorithms are susceptible to spurious local minima, have relatively poor accuracy, and tend to be very slow unless axisymmetry is assumed. When the underlying simulation uses spectral methods, then minimization algorithms can be somewhat faster and more robust.

Spectral integral-iteration algorithms (Section 8.4) and elliptic-PDE algorithms (Section 8.5) are fast and accurate, but are moderately difficult to program. Their main disadvantage is the need for a fairly good initial guess for the horizon position/shape.

In many cases Schnetter's "pretracking" algorithm (Section 8.6) can greatly improve an elliptic-PDE algorithm's robustness, by determining – *before* it appears – approximately where (in space) and when (in time) a new outermost apparent horizon will appear. Pretracking is implemented as a modification of an existing elliptic-PDE algorithm and is moderately slow: It typically has a cost 5 to 10 times that of finding a single horizon with the elliptic-PDE algorithm.

Finally, flow algorithms (Section 8.7) are generally quite slow (Metzger's algorithm [109] is a notable exception) but can be very robust in their convergence. They are moderately easy to program. Flow algorithms are global algorithms, in that their convergence does not depend on having a good initial guess.

Table 2 lists freely-available apparent horizon finding codes.

10 Acknowledgements

I thank the many researchers who answered my e-mail queries on various aspects of their work. I thank the anonymous referees for their careful reading of the manuscript and their many helpful comments. I thank Badri Krishnan for many useful conversations on the properties of apparent, isolated, and dynamical horizons. I thank Scott Caveny and Peter Diener for useful conversations on event-horizon finders. I thank Peter Diener, Luciano Rezzolla, and Virginia J. Vitzthum for helpful comments on various drafts of this paper. I thank Peter Diener and Edward Seidel for providing unpublished figures.

I thank the many authors named in this review for granting permission to reprint figures from their published work. I thank the American Astronomical Society, the American Physical Society, and IOP Publishing for granting permission to reprint figures published in their journals. The American Physical Society requires the following disclaimer regarding such reprinted material:

Readers may view, browse, and/or download material for temporary copying purposes only, provided these uses are for noncommercial personal purposes. Except as provided by law, this material may not be further reproduced, distributed, transmitted, modified, adapted, performed, displayed, published, or sold in whole or part, without written permission from the publisher.

I thank the Alexander von Humboldt Foundation, the AEI visitors program, and the AEI postdoctoral fellowship program for financial support.

A Solving a Single Nonlinear Algebraic Equation

In this appendix I briefly outline numerical algorithms and codes for solving a single 1-dimensional nonlinear algebraic equation $f(x) = 0$, where the continuous function $f : \mathfrak{R} \rightarrow \mathfrak{R}$ is given.

The process generally begins by evaluating f on a suitable grid of points and looking for sign changes. By the intermediate value theorem, each sign change must bracket at least one root. Given a pair of such ordinates x_- and x_+ , there are a variety of algorithms available to accurately and efficiently find the (a) root:

If $|x_+ - x_-|$ is small, say on the order of a finite-difference grid spacing, then closed-form approximations are probably accurate enough:

- The simplest approximation is a simple linear interpolation of f between x_- and x_+ .
- A slightly more sophisticated algorithm, “inverse quadratic interpolation”, is to use three ordinates, two of which bracket a root, and estimate the root as the root of the (unique) parabola which passes through the three given $(x, f(x))$ points⁶⁴.

For larger $|x_+ - x_-|$, iterative algorithms are necessary to obtain an accurate root:

- Bisection (binary search on the sign of f) is a well-known iterative scheme which is very robust. However, it is rather slow if high accuracy is desired.
- Newton’s method can be used, but it requires that the derivative f' be available. Alternatively, the secant algorithm (similar to Newton’s method but estimating f' from the most recent pair of function evaluations) gives similarly fast convergence without requiring f' to be available. Unfortunately, if $|f'|$ is small enough at any iteration point, both these algorithms can fail to converge, or more generally they can generate “wild” trial ordinates.
- Probably the most sophisticated algorithm is that of van Wijngaarden, Dekker, and Brent. This is a carefully engineered hybrid of the bisection, secant, and inverse quadratic interpolation algorithms, and generally combines the rapid convergence of the secant algorithm with the robustness of bisection. The van Wijngaarden–Dekker–Brent algorithm is described by Forsythe, Malcolm, and Moler [71, Chapter 7], Kahaner, Moler, and Nash [92, Chapter 7], and Press et al. [125, Section 9.3]. An excellent implementation of this, the Fortran subroutine ZEROIN, is freely available from <http://www.netlib.org/fmm/>.

⁶⁴The parabola generically has two roots, but normally only one of them lies between x_- and x_+ .

B The Numerical Integration of Ordinary Differential Equations

The time-integration problem⁶⁵ for ordinary differential equations (ODEs) is traditionally written as follows: We are given an integer $n > 0$ (the number of ODEs to integrate), a “right-hand-side” function $f : \mathfrak{R}^n \times \mathfrak{R} \rightarrow \mathfrak{R}^n$, and the value $y(0)$ of a function $y : \mathfrak{R} \rightarrow \mathfrak{R}^n$ satisfying the ODEs

$$\frac{dy}{dt} = f(y, t) \quad (42)$$

We wish to know (or approximate) $y(t)$ for some finite interval $t \in [0, t_{\max}]$.

This is a well-studied problem in numerical analysis. See, for example, Forsythe, Malcolm, and Moler [71, Chapter 6] or Kahaner, Moler, and Nash [92, Chapter 8] for a general overview of ODE integration algorithms and codes, or Shampine and Gordon [140], Hindmarsh [84], or Brankin, Gladwell, and Shampine [38] for detailed technical accounts.

For our purposes, it suffices to note that highly accurate, efficient, and robust ODE-integration codes are widely available. In fact, there is a strong tradition in numerical analysis of free availability of such codes. Notably, Table 3 lists several freely-available ODE codes. As well as being of excellent numerical quality, these codes are also very easy to use, employing sophisticated adaptive algorithms to automatically adjust the step size and/or the precise integration scheme used⁶⁶. These codes can generally be relied upon to produce accurate results both more efficiently and more easily than a hand-crafted integrator. I have used the LSODE solver in several research projects with excellent results.

Program	Reference(s)	Web page
RKF45	[71, Chapter 6]	http://www.netlib.org/ode/rkf45.f
ODE (DE/STEP)	[140]	http://www.netlib.org/ode/ode.f
ODEPACK (LSODE/LSODA etc.)	[84]	http://www.netlib.org/odepack/
RKSUITE	[38]	http://www.netlib.org/ode/rksuite/

Table 3: *This table lists some general-purpose ODE codes which are freely available.*

⁶⁵The numerical-analysis literature usually refers to this as the “initial value problem”. Unfortunately, in a relativity context this terminology often causes confusion with the “initial data problem” of solving the ADM constraint equations. I use the term “time-integration problem for ODEs” to (try to) avoid this confusion. In this appendix, sans-serif lower-case letters $abc\dots z$ denote variables and functions in \mathfrak{R}^n (for some fixed dimension n), and sans-serif upper-case letters $ABC\dots Z$ denote $n \times n$ real-valued matrices.

⁶⁶LSODA can also automatically detect stiff systems of ODEs and adjust its integration scheme so as to handle them efficiently.

References

- [1] Abrahams, A.M., Cook, G.B., Shapiro, S.L., and Teukolsky, S.A., “Solving Einstein’s Equations for Rotating Spacetimes: Evolution of Relativistic Star Clusters”, *Phys. Rev. D*, **49**, 5153–5164, (1994). 4
- [2] Abrahams, A.M., and Evans, C.R., “Trapping a Geon: Black Hole Formation by an Imploding Gravitational Wave”, *Phys. Rev. D*, **46**, R4117–R4121, (1992). 8.2
- [3] Abrahams, A.M., Heiderich, K.H., Shapiro, S.L., and Teukolsky, S.A., “Vacuum initial data, singularities, and cosmic censorship”, *Phys. Rev. D*, **46**, 2452–2463, (1992). 8.2
- [4] Alcubierre, M., Brandt, S., Brügmann, B., Gundlach, C., Massó, J., Seidel, E., and Walker, P., “Test-beds and applications for apparent horizon finders in numerical relativity”, *Class. Quantum Grav.*, **17**, 2159–2190, (2000). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9809004>. 3.2, 8.2, 8.3.4, 8.7.4, 8.7.5
- [5] Alcubierre, M., Brügmann, B., Diener, P., Guzmán, F.S., Hawke, I., Hawley, S., Herrmann, F., Koppitz, M., Pollney, D., Seidel, E., and Thornburg, J., “Dynamical evolution of quasi-circular binary black hole data”, *Phys. Rev. D*, **72**, 044004, (2005). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0411149>. 3.3, 11, 8.7.5
- [6] Andersson, L., and Metzger, J., personal communication, (2007). Personal communication from Lars Andersson to Bela Szilágyi. 17
- [7] Anninos, P., Bernstein, D., Brandt, S., Libson, J., Massó, J., Seidel, E., Smarr, L.L., Suen, W.-M., and Walker, P., “Dynamics of Apparent and Event Horizons”, *Phys. Rev. Lett.*, **74**, 630–633, (1995). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9403011>. 5.2, 5.3, 5.3.2, 13
- [8] Anninos, P., Camarda, K., Libson, J., Massó, J., Seidel, E., and Suen, W.-M., “Finding apparent horizons in dynamic 3D numerical spacetimes”, *Phys. Rev. D*, **58**, 024003, 1–12, (1998). 8.3, 8.3.1, 8.3.4
- [9] Anninos, P., Daues, G., Massó, J., Seidel, E., and Suen, W.-M., “Horizon boundary conditions for black hole spacetimes”, *Phys. Rev. D*, **51**, 5562–5578, (1995). 8.1
- [10] Ansorg, M., “A double-domain spectral method for black hole excision data”, *Phys. Rev. D*, **72**, 024018, 1–10, (2005). Related online version (cited on 30 January 2007): <http://arxiv.org/abs/gr-qc/0505059>. 26, 38
- [11] Ansorg, M., Brügmann, B., and Tichy, W., “Single-domain spectral method for black hole puncture data”, *Phys. Rev. D*, **70**, 064011, 1–13, (2004). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0404056>. 26, 38
- [12] Ansorg, M., Kleinwächter, A., and Meinel, R., “Highly accurate calculation of rotating neutron stars: Detailed description of the numerical methods”, *Astron. Astrophys.*, **405**, 711–721, (2003). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/astro-ph/0301173>. 26, 38
- [13] Ansorg, M., and Petroff, D., “Black holes surrounded by uniformly rotating rings”, *Phys. Rev. D*, **72**, 024019, (2005). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0505060>. 26, 38

- [14] Arnowitt, R., Deser, S., and Misner, C.W., “The dynamics of general relativity”, in Witten, L., ed., *Gravitation: An Introduction to Current Research*, 227–265, (Wiley, New York, U.S.A., 1962). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0405109>. I, 1
- [15] Ascher, U.M., Mattheij, R.M.M., and Russell, R.D., *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations*, (Prentice-Hall, Englewood Cliffs, U.S.A., 1988). 31
- [16] Ashtekar, A., Beetle, C., and Fairhurst, S., “Isolated horizons: a generalization of black hole mechanics”, *Class. Quantum Grav.*, **16**, L1–L7, (1999). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/9812065>. 7.3
- [17] Ashtekar, A., and Galloway, G., “Some uniqueness results for dynamical horizons”, *Adv. Theor. Math. Phys.*, **9**, 1–30, (2005). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0503109>. 22
- [18] Ashtekar, A., and Krishnan, B., “Dynamical Horizons: Energy, Angular Momentum, Fluxes, and Balance Laws”, *Phys. Rev. Lett.*, **89**, 261101, 1–4, (2002). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0207080>. 7.3
- [19] Ashtekar, A., and Krishnan, B., “Dynamical horizons and their properties”, *Phys. Rev. D*, **68**, 104030, 1–25, (2003). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0308033>. 7.3
- [20] Ashtekar, A., and Krishnan, B., “Isolated and Dynamical Horizons and Their Applications”, *Living Rev. Relativity*, **7**, lrr-2004-10, 10, (2004). URL (cited on 09 January 2006):
<http://www.livingreviews.org/lrr-2004-10>. 7.3
- [21] Baiotti, L., Hawke, I., Montero, P.J., Löffler, F., Rezzolla, L., Stergioulas, N., Font, J.A., and Seidel, E., “Three-dimensional relativistic simulations of rotating neutron star collapse to a Kerr black hole”, *Phys. Rev. D*, **71**, 024035, (2005). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0403029>. 8
- [22] Balay, S., Buschelman, K., Gropp, W.D., Kaushik, D., Knepley, M., Curfman McInnes, L., Smith, B.F., and Zhang, H., “PETSc: Portable, Extensible Toolkit for Scientific Computation”, project homepage, Argonne National Laboratory. URL (cited on 09 January 2006):
<http://www.mcs.anl.gov/petsc>. 8.5.3
- [23] Balay, S., Buschelman, K., Gropp, W.D., Kaushik, D., Knepley, M., Curfman McInnes, L., Smith, B.F., and Zhang, H., *PETSc Users Manual*, ANL-95/11 – Revision 2.1.5, (Argonne National Laboratory, Argonne, U.S.A., 2003). URL (cited on 20 August 2003):
<http://www-unix.mcs.anl.gov/petsc/petsc-as/documentation/>. 8.5.3
- [24] Balay, S., Gropp, W.D., Curfman McInnes, L., and Smith, B.F., “Efficient Management of Parallelism in Object-Oriented Numerical Software Libraries”, in Arge, E., Bruaset, A.M., and Langtangen, H.P., eds., *Modern Software Tools for Scientific Computing*, Proceedings of SciTools ’96 Workshop held in Oslo, Norway, 163–202, (Birkhäuser, Boston, U.S.A., 1997). 8.5.3
- [25] Bartnik, R., personal communication. Personal communication from Robert Bartnik to Carsten Gundlach. 8.7

- [26] Baumgarte, T.W., Cook, G.B., Scheel, M.A., Shapiro, S.L., and Teukolsky, S.A., “Implementing an apparent-horizon finder in three dimensions”, *Phys. Rev. D*, **54**, 4849–4857, (1996). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9606010>. 8.3, 8.3.1, 8.3.4
- [27] Baumgarte, T.W., and Shapiro, S.L., “Numerical relativity and compact binaries”, *Phys. Rep.*, **376**, 41–131, (2003). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0211028>. I, 23
- [28] Bernstein, D., *Notes on the Mean Curvature Flow Method for Finding Apparent Horizons*, (National Center for Supercomputing Applications, Urbana-Champaign, U.S.A., 1993). 8.7, 8.7.1
- [29] Bishop, N.T., “The Closed Trapped Region and the Apparent Horizon of Two Schwarzschild Black Holes”, *Gen. Relativ. Gravit.*, **14**, 717–723, (1982). 8.2
- [30] Bishop, N.T., “The horizons of two Schwarzschild black holes”, *Gen. Relativ. Gravit.*, **16**, 589–593, (1984). 8.2
- [31] Bishop, N.T., “The Event Horizons of Two Schwarzschild black holes”, *Gen. Relativ. Gravit.*, **20**, 573–581, (1988). 5.1
- [32] Bizoń, P., Malec, E., and Ó Murchadha, N., “Trapped Surfaces in Spherical Stars”, *Phys. Rev. Lett.*, **61**, 1147–1150, (1988). 30
- [33] Bonazzola, S., Friebe, J., Gourgoulhon, E., and Marck, J.-A., “Spectral methods in general relativity – toward the simulation of 3D-gravitational collapse of neutron stars”, in Ilin, A.V., and Scott, L.R., eds., *ICOSAHOM '95*, Proceedings of the Third International Conference on Spectral and High Order Methods: Houston, Texas, June 5–9, 1995, Houston Journal of Mathematics, 3–19, (University of Houston, Houston, U.S.A., 1996). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9604029>. 26, 38
- [34] Bonazzola, S., Gourgoulhon, E., and Marck, J.-A., “Spectral methods in general relativistic astrophysics”, *J. Comput. Appl. Math.*, **109**, 433–473, (1999). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9811089>. 26, 38
- [35] Bonazzola, S., and Marck, J.-A., “Pseudo-Spectral Methods Applied to Gravitational Collapse”, in Evans, C.R., Finn, L.S., and Hobill, D.W., eds., *Frontiers in Numerical Relativity*, Proceedings of the International Workshop on Numerical Relativity, University of Illinois at Urbana-Champaign, USA, May 9–13, 1988, 239–253, (Cambridge University Press, Cambridge, U.K.; New York, U.S.A., 1989). 26, 38
- [36] Booth, I., “Black hole boundaries”, *Can. J. Phys.*, **83**, 1073–1099, (2005). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0508107>. 7.3
- [37] Boyd, J.P., *Chebyshev and Fourier Spectral Methods*, (Dover Publications, Mineola, U.S.A., 2001), 2nd edition. 26, 38, 8.5.3
- [38] Brankin, R.W., Gladwell, I., and Shampine, L.F., “RKSUITE: A Suite of Runge–Kutta Codes for the Initial Value Problem for ODEs”, other, Dept. of Mathematics, Southern Methodist University, Dallas, TX, (1992). URL (cited on 09 January 2006): <http://www.netlib.org/ode/rksuite/>. B

- [39] Brent, R.P., *Algorithms for Minimization Without Derivatives*, (Dover Publications, Mineola, U.S.A., 2002). Reprint of 1973 original edition. 32
- [40] Brewin, L.C., “Is the Regge Calculus a Consistent Approximation to General Relativity?”, *Gen. Relativ. Gravit.*, **32**, 897–918, (2000). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9502043>. 4
- [41] Brewin, L.C., and Gentle, A.P., “On the Convergence of Regge Calculus to General Relativity”, *Class. Quantum Grav.*, **18**, 517–525, (2001). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0006017>. 4
- [42] Briggs, W.L., Henson, V.E., and McCormick, S.F., *A Multigrid Tutorial*, (SIAM, Philadelphia, U.S.A., 2000), 2nd edition. 44, 62
- [43] Brill, D.R., and Lindquist, R.W., “Interaction Energy in Geometrostatics”, *Phys. Rev.*, **131**, 471–476, (1963). 8.3.4
- [44] Caveny, S.A., *Tracking Black Holes in Numerical Relativity: Foundations and Applications*, Ph.D. Thesis, (University of Texas at Austin, Austin, U.S.A., 2002). 5.3.2, 5.3.3
- [45] Caveny, S.A., Anderson, M., and Matzner, R.A., “Tracking Black Holes in Numerical Relativity”, *Phys. Rev. D*, **68**, 104009, (2003). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0303099>. 5.3.3
- [46] Caveny, S.A., and Matzner, R.A., “Adaptive event horizon tracking and critical phenomena in binary black hole coalescence”, *Phys. Rev. D*, **68**, 104003–1–13, (2003). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0303109>. 5.3.2, 6
- [47] Choptuik, M.W., *A Study of Numerical Techniques for Radiative Problems in General Relativity*, Ph.D. Thesis, (University of British Columbia, Vancouver, Canada, 1986). 8.1
- [48] Choptuik, M.W., “Experiences with an Adaptive Mesh Refinement Algorithm in Numerical Relativity”, in Evans, C.R., Finn, L.S., and Hobill, D.W., eds., *Frontiers in Numerical Relativity*, Proceedings of the International Workshop on Numerical Relativity, University of Illinois at Urbana-Champaign (Urbana-Champaign, Illinois, USA), May 9–13, 1988, 206–221, (Cambridge University Press, Cambridge, U.K.; New York, U.S.A., 1989). 5
- [49] Chruściel, P.T., and Galloway, G.J., “Horizons Non-Differentiable on a Dense Set”, *Commun. Math. Phys.*, **193**, 449–470, (1998). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9611032>. 6
- [50] Cook, G.B., *Initial Data for the Two-Body Problem of General Relativity*, Ph.D. Thesis, (University of North Carolina at Chapel Hill, Chapel Hill, U.S.A., 1990). 8.5.7
- [51] Cook, G.B., and Abrahams, A.M., “Horizon Structure of Initial-Data Sets for Axisymmetric Two-Black-Hole Collisions”, *Phys. Rev. D*, **46**, 702–713, (1992). 7.5, 8.5.7
- [52] Cook, G.B., and York Jr, J.W., “Apparent Horizons for Boosted or Spinning Black Holes”, *Phys. Rev. D*, **41**, 1077–1085, (1990). 8.5.7
- [53] Curtis, A.R., and Reid, J.K., “The Choice of Step Lengths When Using Differences to Approximate Jacobian Matrices”, *J. Inst. Math. Appl.*, **13**, 121–126, (1974). 8.5.4

- [54] Davis, T.A., “UMFPACK: unsymmetric multifrontal sparse LU factorization package”, project homepage, University of Florida (CISE). URL (cited on 6 January 2007): <http://www.cise.ufl.edu/research/sparse/umfpack/>. 8.5.5
- [55] Davis, T.A., “Algorithm 832: UMFPACK V4.3—an unsymmetric-pattern multifrontal method”, *ACM Trans. Math. Software*, **30**, 196–199, (2004). Related online version (cited on 09 January 2006): <http://www.cise.ufl.edu/~davis/>. TR-02-002. 8.5.5
- [56] Davis, T.A., “A column pre-ordering strategy for the unsymmetric-pattern multifrontal method”, *ACM Trans. Math. Software*, **30**, 165–195, (2004). Related online version (cited on 09 January 2006): <http://www.cise.ufl.edu/~davis/>. TR-02-001. 8.5.5
- [57] Davis, T.A., and Duff, I.S., “An unsymmetric-pattern multifrontal method for sparse LU factorization”, *SIAM J. Matrix Anal. Appl.*, **18**, 140–158, (1997). 8.5.5
- [58] Davis, T.A., and Duff, I.S., “A combined unifrontal/multifrontal method for unsymmetric sparse matrices”, *ACM Trans. Math. Software*, **25**, 1–19, (1999). 8.5.5
- [59] Dennis Jr, J.E., and Schnabel, R.B., *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, (SIAM, Philadelphia, U.S.A., 1996). 32
- [60] Diener, P., “A New General Purpose Event Horizon Finder for 3D”, *Class. Quantum Grav.*, **20**, 4901–4917, (2003). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0305039>. 3.2, 11, 5.3.3, 5.3.3, 5.3.3, 14, 6
- [61] Diener, P., personal communication, (2007). 5.3.3
- [62] Diener, P., Herrmann, F., Pollney, D., Schnetter, E., Seidel, E., Takahashi, R., Thornburg, J., and Ventrella, J., “Accurate Evolution of Orbiting Binary Black Holes”, *Phys. Rev. Lett.*, **96**, 121101, (2006). Related online version (cited on 3 October 2006): <http://arXiv.org/abs/gr-qc/0512108>. 8.7.5
- [63] Dreyer, O., Krishnan, B., Schnetter, E., and Shoemaker, D., “Introduction to isolated horizons in numerical relativity”, *Phys. Rev. D*, **67**, 024018, 1–14, (2003). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0206008>. 7.3
- [64] Du Fort, E.C., and Frankel, S.P., “Stability Conditions in the Numerical Treatment of Parabolic Differential Equations”, *Math. Tables Aids Comput.*, **7**, 135–152, (1953). 8.7.1
- [65] Duff, I.S., Erisman, A.M., and Reid, J.K., *Direct Methods for Sparse Matrices*, (Oxford University Press, Oxford, U.K.; New York, U.S.A., 1986). 8.5.5
- [66] Dykema, P.G., *The Numerical Simulation of Axially Symmetric Gravitational Collapse*, Ph.D. Thesis, (University of Texas at Austin, Austin, U.S.A., 1980). 8.2
- [67] Eardley, D.M., “Gravitational Collapse of Marginally Bound Spheroids: Initial Conditions”, *Phys. Rev. D*, **12**, 3072–3076, (1975). 8.5.7
- [68] Eppley, K.R., *The numerical evolution of the collision of two black holes*, Ph.D. Thesis, (Princeton University, Princeton, U.S.A., 1975). 47
- [69] Eppley, K.R., “Evolution of time-symmetric gravitational waves: Initial data and apparent horizons”, *Phys. Rev. D*, **16**, 1609–1614, (1977). 8.3.4

- [70] Fornberg, B., *A Practical Guide to Pseudospectral Methods*, Cambridge Monographs on Applied and Computational Mathematics, (Cambridge University Press, Cambridge, U.K.; New York, U.S.A., 1998). 26, 38
- [71] Forsythe, G.E., Malcolm, M.A., and Moler, C.B., *Computer Methods for Mathematical Computations*, (Prentice-Hall, Englewood Cliffs, U.S.A., 1977). Related online version (cited on 09 January 2006):
<http://www.netlib.org/fmm/>. A, B
- [72] Gentle, A.P., “Regge Calculus: A Unique Tool for Numerical Relativity”, *Gen. Relativ. Gravit.*, **34**, 1701–1718, (2002). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0408006>. 2.3
- [73] Gentle, A.P., and Miller, W.A., “A fully (3+1)-dimensional Regge calculus model of the Kasner cosmology”, *Class. Quantum Grav.*, **15**, 389–405, (1998). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/9706034>. 4
- [74] Goodale, T., Allen, G., Lanfermann, G., Massó, J., Radke, T., Seidel, E., and Shalf, J., “The Cactus Framework and Toolkit: Design and Applications”, in Palma, J.M.L.M. and Dongarra, J., Hernández, V., and Sousa, A.A., eds., *High Performance Computing for Computational Science (VECPAR 2002)*, 5th International Conference, Porto, Portugal, June 26–28, 2002: Selected papers and invited talks, vol. 2565 of Lecture Notes in Computer Science, 197–227, (Springer, Berlin, Germany; New York, U.S.A., 2003). 3.1
- [75] Gottlieb, D., and Orszag, S.A., *Numerical Analysis of Spectral Methods: Theory and Applications*, vol. 26 of Regional Conference Series in Applied Mathematics, (SIAM, Philadelphia, U.S.A., 1977). Based on a series of lectures presented at the NSF-CBMS regional conference held at Old Dominion University from August 2–6, 1976. 26, 38
- [76]ourgoulhon, E., and Jaramillo, J.L., “A 3+1 perspective on null hypersurfaces and isolated horizons”, *Phys. Rep.*, **423**, 159–294, (2006). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0503113>. 7.3
- [77] Grandclément, P., Bonazzola, S.,ourgoulhon, E., and Marck, J.-A., “A Multidomain Spectral Method for Scalar and Vectorial Poisson Equations with Noncompact Sources”, *J. Comput. Phys.*, **170**, 231–260, (2001). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0003072>. 26, 38
- [78] Grandclément, P.,ourgoulhon, E., and Bonazzola, S., “Binary black holes in circular orbits. II. Numerical methods and first results”, *Phys. Rev. D*, **65**, 044021, 1–18, (2002). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0106016>. 7
- [79] Grayson, M.A., “The Heat Equation Shrinks Embedded Plane Curves to Round Points”, *J. Differ. Geom.*, **26**, 285–314, (1987). 8.7
- [80] Gundlach, C., “Pseudo-spectral apparent horizon finders: An efficient new algorithm”, *Phys. Rev. D*, **57**, 863–875, (1998). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/9707050>. 23, 8.4, 56, 8.7.4, 8.7.5
- [81] Hawking, S.W., “The Event Horizon”, in DeWitt, C., and DeWitt, B.S., eds., *Black Holes*, Based on lectures given at the 23rd session of the Summer School of Les Houches, 1972, 1–56, (Gordon and Breach, New York, U.S.A., 1973). 4, 18

- [82] Hawking, S.W., and Ellis, G.F.R., *The Large Scale Structure of Space-Time*, Cambridge Monographs on Mathematical Physics, (Cambridge University Press, Cambridge, U.K., 1973). 4, 7.1, 19
- [83] Hayward, S.A., “General laws of black hole dynamics”, *Phys. Rev. D*, **49**, 6467–6474, (1994). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9306006>. 7.3
- [84] Hindmarsh, A.C., “ODEPACK, A Systematized Collection of ODE Solvers”, in Stepleman, R.S. et al., ed., *Scientific Computing: Applications of Mathematics and Computing to the Physical Sciences*, Based on papers presented at the Tenth IMACS World Congress on System Simulation and Scientific Computation, held in Montreal, Canada, August 8–13, 1982, vol. 1 of IMACS Transactions on Scientific Computing, 55–64, (North-Holland, Amsterdam, Netherlands; New York, U.S.A., 1983). Related online version (cited on 09 January 2006): <http://www.netlib.org/odepack/index.html>. B
- [85] Hochbruck, M., Lubich, C., and Selhofer, H., “Exponential Integrators for Large Systems of Differential Equations”, *SIAM J. Sci. Comput.*, **19**, 1552–1574, (1998). 8.7.1
- [86] Hornung, R.D., and Kohn, S.R., “Managing application complexity in the SAMRAI object-oriented framework”, *Concurr. Comput. Pract. Exp.*, **14**, 347–368, (2002). 3.1
- [87] Hornung, R.D., Wissink, A.M., and Kohn, S.R., “Managing complex data and geometry in parallel structured AMR applications”, *Eng. Comput.*, **22**, 181–195, (2006). 3.1
- [88] Hughes, S.A., Keeton II, C.R., Walker, P., Walsh, K.T., Shapiro, S.L., and Teukolsky, S.A., “Finding Black Holes in Numerical Spacetimes”, *Phys. Rev. D*, **49**, 4004–4015, (1994). 5.1, 5.1.2
- [89] Huq, M.F., *Apparent Horizon Location in Numerical Spacetimes*, Ph.D. Thesis, (The University of Texas at Austin, Austin, U.S.A., 1996). 25, 8.5.1, 8.5.2, 2, 8.5.3, 8.5.4, 8.5.7
- [90] Huq, M.F., Choptuik, M.W., and Matzner, R.A., “Locating Boosted Kerr and Schwarzschild Apparent Horizons”, *Phys. Rev. D*, **66**, 084024, (2002). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0002076>. 25, 8.5.1, 8.5.2, 2, 8.5.3, 8.5.4, 8.5.7
- [91] Husa, S., and Winicour, J., “Asymmetric merger of black holes”, *Phys. Rev. D*, **60**, 084019, 1–13, (1999). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9905039>. 4
- [92] Kahaner, D., Moler, C.B., and Nash, S., *Numerical Methods and Software*, (Prentice Hall, Englewood Cliffs, U.S.A., 1989). Revised and (greatly) expanded edition of Forsythe, G.E. and Malcolm, M.A. and Moler, C.B., “Computer methods for mathematical computations”(1977). A, B
- [93] Kembell, A.J., and Bishop, N.T., “The numerical determination of apparent horizons”, *Class. Quantum Grav.*, **8**, 1361–1367, (1991). 8.4.1, 8.4.3
- [94] Kershaw, D.S., “The Incomplete Cholesky-Conjugate Gradient Method for Iterative Solution of Linear Equations”, *J. Comput. Phys.*, **26**, 43–65, (1978). 8.5.5, 45
- [95] Kidder, L.E., and Finn, L.S., “Spectral Methods for Numerical Relativity. The Initial Data Problem”, *Phys. Rev. D*, **62**, 084026, (2000). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9911014>. 26, 38

- [96] Kidder, L.E., Scheel, M.A., Teukolsky, S.A., Carlson, E.D., and Cook, G.B., “Black hole evolution by spectral methods”, *Phys. Rev. D*, **62**, 084032, (2000). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0005056>. 26, 38
- [97] Kidder, L.E., Scheel, M.A., Teukolsky, S.A., and Cook, G.B., “Spectral Evolution of Einstein’s Equations”, Miniprogram on Colliding Black Holes: Mathematical Issues in Numerical Relativity, held at the Institute for Theoretical Physics, UC at Santa Barbara, 10–28 January 2000, conference paper, (2000). 26, 38
- [98] Kriele, M., and Hayward, S.A., “Outer trapped surfaces and their apparent horizon”, *J. Math. Phys.*, **38**, 1593–1604, (1997). 7.1
- [99] Lehner, L., Bishop, N.T., Gómez, R., Szilágyi, B., and Winicour, J., “Exact solutions for the intrinsic geometry of black hole coalescence”, *Phys. Rev. D*, **60**, 044005, 1–10, (1999). 4
- [100] Lehner, L., Gómez, R., Husa, S., Szilágyi, B., Bishop, N.T., and Winicour, J., “Bagels Form When Black Holes Collide”, institutional homepage, Pittsburgh Supercomputing Center. URL (cited on 09 January 2006):
<http://www.psc.edu/research/graphics/gallery/winicour.html>. 4
- [101] Leiler, G., and Rezzolla, L., “On the iterated Crank–Nicolson method for hyperbolic and parabolic equations in numerical relativity”, *Phys. Rev. D*, **73**, 044001, (2006). Related online version (cited on 3 October 2006):
<http://arXiv.org/abs/gr-qc/0601139>. 61
- [102] Libson, J., Massó, J., Seidel, E., and Suen, W.-M., “A 3D Apparent Horizon Finder”, in Jantzen, R.T., and Keiser, G.M., eds., *The Seventh Marcel Grossmann Meeting: On recent developments in theoretical and experimental general relativity, gravitation, and relativistic field theories*, Proceedings of the meeting held at Stanford University, July 24–30, 1994, 631, (World Scientific, Singapore; River Edge, U.S.A., 1996). 8.3.4
- [103] Libson, J., Massó, J., Seidel, E., Suen, W.-M., and Walker, P., “Event horizons in numerical relativity: Methods and tests”, *Phys. Rev. D*, **53**, 4335–4350, (1996). 5.2, 5.3, 5.3.2, 13, 4
- [104] Lin, L.-M., and Novak, J., “Three-dimensional apparent horizon finder in LORENE”, personal communication, (2006). Personal communication from Lap-Ming Lin to Jonathan Thornburg. 3.2, 8.4.2, 8.4.2, 8.4.2, 8.4.3
- [105] Lorensen, W.E., and Cline, H.E., “Marching cubes: A high resolution 3D surface construction algorithm”, *SIGGRAPH Comput. Graph.*, **21**, 163–169, (1987). 2.1
- [106] MacNeice, P., Olson, K.M., Mobarry, C., de Fainchtein, R., and Packer, C., “PARAMESH: A parallel adaptive mesh refinement community toolkit”, *Computer Phys. Commun.*, **126**, 330–354, (2000). 3.1
- [107] Madderom, P., “Incomplete LU-Decomposition – Conjugate Gradient”, unknown format, (1984). Fortran 66 subroutine. 45
- [108] Matzner, R.A., Seidel, E., Shapiro, S.L., Smarr, L.L., Suen, W.-M., Teukolsky, S.A., and Winicour, J., “Geometry of a Black Hole Collision”, *Science*, **270**, 941–947, (1995). 5, 8.5.6
- [109] Metzger, J., “Numerical computation of constant mean curvature surfaces using finite elements”, *Class. Quantum Grav.*, **21**, 4625–4646, (2004). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0408059>. 8.7.1, 8.7.3, 8.7.5, 9

- [110] Miller, M.A., “Regge Calculus as a Fourth Order Method in Numerical Relativity”, *Class. Quantum Grav.*, **12**, 3037–3051, (1995). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9502044>. 4
- [111] Misner, C.W., and Sharp, D.H., “Relativistic Equations for Adiabatic, Spherically Symmetric Gravitational Collapse”, *Phys. Rev.*, **136**, B571–B576, (1964). 9
- [112] Misner, C.W., Thorne, K.S., and Wheeler, J.A., *Gravitation*, (W.H. Freeman, San Francisco, U.S.A., 1973). 1, 9
- [113] Nakamura, T., Kojima, Y., and Oohara, K., “A Method of Determining Apparent Horizons in Three-Dimensional Numerical Relativity”, *Phys. Lett. A*, **106**, 235–238, (1984). 8.4, 8.4, 8.4
- [114] Oohara, K., “Apparent Horizon of Initial Data for Black Hole-Collisions”, in Sato, H., and Nakamura, T., eds., *Gravitational Collapse and Relativity*, Proceedings of Yamada Conference XIV, Kyoto International Conference Hall, Japan, April 7–11, 1986, 313–319, (World Scientific, Singapore; Philadelphia, U.S.A., 1986). 8.4.3
- [115] Oohara, K., Nakamura, T., and Kojima, Y., “Apparent Horizons of Time-Symmetric Initial Value for Three Black Holes”, *Phys. Lett. A*, **107**, 452–455, (1985). 8.4.3
- [116] Osher, S., and Sethian, J.A., “Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulations”, *J. Comput. Phys.*, **79**, 12–49, (1988). 8.7.1
- [117] Parashar, M., and Browne, J.C., “System Engineering for High Performance Computing Software: The HDDA/DAGH Infrastructure for Implementation of Parallel Structured Adaptive Mesh Refinement”, in Baden, S.B., Chrisochoides, N.P., Gannon, D.B., and Norman, M.L., eds., *Structured Adaptive Mesh Refinement (SAMR) Grid Methods*, vol. 117 of IMA Volumes in Mathematics and its Applications, 1–18, (Springer, New York, U.S.A., 2000). 3.1
- [118] Pasch, E., *The level set method for the mean curvature flow on (R^3, g)* , SFB 382 Reports, 63, (University of Tübingen, Tübingen, Germany, 1997). URL (cited on 09 January 2006): <http://www.uni-tuebingen.de/uni/opx/reports.html>. 8.5.4, 8.7, 8.7.1, 8.7.3, 60
- [119] Petrich, L.I., Shapiro, S.L., and Teukolsky, S.A., “Oppenheimer–Snyder Collapse with Maximal Time Slicing and Isotropic Coordinates”, *Phys. Rev. D*, **31**, 2459–2469, (1985). 8.1
- [120] Pfeiffer, H.P., *Initial Data for Black Hole Evolutions*, Ph.D. Thesis, (Cornell University, Ithaca, U.S.A., 2003). Related online version (cited on 1 October 2006): <http://arXiv.org/abs/gr-qc/0510016>. 26, 38
- [121] Pfeiffer, H.P., personal communication, (2006). 8.3.1, 35
- [122] Pfeiffer, H.P., Cook, G.B., and Teukolsky, S.A., “Comparing initial-data sets for binary black holes”, *Phys. Rev. D*, **66**, 024047, 1–17, (2002). Related online version (cited on 1 October 2006): <http://arXiv.org/abs/gr-qc/0203085>. 26, 35, 38
- [123] Pfeiffer, H.P., Kidder, L.E., Scheel, M.A., and Teukolsky, S.A., “A multidomain spectral method for solving elliptic equations”, *Computer Phys. Commun.*, **152**, 253–273, (2003). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0202096>. 26, 38

- [124] Pfeiffer, H.P., Teukolsky, S.A., and Cook, G.B., “Quasicircular orbits for spinning binary black holes”, *Phys. Rev. D*, **62**, 104018, 1–11, (2000). Related online version (cited on 1 October 2006):
<http://arXiv.org/abs/gr-qc/0006084>. 7.5, 26, 8.3.1, 38
- [125] Press, W.H., Flannery, B.P., Teukolsky, S.A., and Vetterling, W.T., *Numerical Recipes*, (Cambridge University Press, Cambridge, U.K.; New York, U.S.A., 1992), 2nd edition. 33, 34, A
- [126] Pretorius, F., and Choptuik, M.W., “Adaptive Mesh Refinement for Coupled Elliptic-Hyperbolic Systems”, *J. Comput. Phys.*, **218**, 246–274, (2006). Related online version (cited on 3 October 2006):
<http://arXiv.org/abs/gr-qc/0508110>. 5
- [127] Pretorius, F., and Lehner, L., “Adaptive mesh refinement for characteristic codes”, *J. Comput. Phys.*, **198**, 10–34, (2004). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0302003>. 5
- [128] Regge, T., “General Relativity without Coordinates”, *Nuovo Cimento A*, **19**, 558–571, (1961). 2.3
- [129] Richtmyer, R.D., and Morton, K.W., *Difference Methods for Initial-Value Problems*, (Krieger, Malabar, U.S.A., 1994), 2nd edition. Reprinted second edition of 1967. 59
- [130] Saad, Y., *Iterative Methods for Sparse Linear Systems*, (SIAM, Philadelphia, U.S.A., 2003), 2nd edition. 8.5.5
- [131] Schnetter, E., “CarpetCode: A mesh refinement driver for Cactus”, project homepage, Center for Computation and Technology, Louisiana State University. URL (cited on 09 January 2006):
<http://www.carpetcode.org>. 3.1, 5.3.3
- [132] Schnetter, E., “A fast apparent horizon algorithm”, (2002). URL (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0206003>. 3.2, 8.5.1, 8.5.2, 39, 8.5.3, 8.5.4, 8.5.5, 8.5.7
- [133] Schnetter, E., “Finding Apparent Horizons and other Two-Surfaces of Constant Expansion”, *Class. Quantum Grav.*, **20**, 4719–4737, (2003). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0306006>. 3.2, 8.5.1, 8.5.2, 39, 8.5.3, 8.5.4, 8.5.5, 8.5.7, 8.6, 8.6.1, 8.6.2, 8.6.4
- [134] Schnetter, E., Hawley, S.H., and Hawke, I., “Evolutions in 3D numerical relativity using fixed mesh refinement”, *Class. Quantum Grav.*, **21**, 1465–1488, (2004). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0310042>. 3.1, 5, 5.3.3
- [135] Schnetter, E., Herrmann, F., and Pollney, D., “Horizon Pretracking”, *Phys. Rev. D*, **71**, 044033, (2005). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0410081>. 3.2, 8.5.7, 8.6, 8.6.1, 8.6.2, 8.6.2, 8.6.3, 51, 8.6.4, 8.6.4, 13, 8.6.6, 14
- [136] Schnetter, E., and Krishnan, B., “Nonsymmetric trapped surfaces in the Schwarzschild Vaidya spacetimes”, *Phys. Rev. D*, **73**, 021502, (2006). Related online version (cited on 3 October 2006):
<http://arXiv.org/abs/gr-qc/0511017>. 20

- [137] Schnetter, E., Krishnan, B., and Beyer, F., “Introduction to Dynamical Horizons in numerical relativity”, *Phys. Rev. D*, **74**, 024028, (2006). Related online version (cited on 30 October 2006):
<http://arXiv.org/abs/gr-qc/0604015>. 7.3
- [138] Schroeder, M.R., *Number Theory in Science and Communication: With Applications in Cryptography, Physics, Digital Information, Computing and Self-Similarity*, vol. 7 of Springer Series in Information Sciences, (Springer, Berlin, Germany; New York, U.S.A., 1986), 2nd edition. 2.2
- [139] Seidel, E., and Suen, W.-M., “Towards a Singularity-Proof Scheme in Numerical Relativity”, *Phys. Rev. Lett.*, **69**, 1845–1848, (1992). 8.1
- [140] Shampine, L.F., and Gordon, M.K., *Computer solution of Ordinary Differential Equations*, (W.H. Freeman, San Francisco, U.S.A., 1975). B
- [141] Shapiro, S.L., and Teukolsky, S.A., “Gravitational Collapse of Supermassive Stars to Black Holes: Numerical Solution of the Einstein Equations”, *Astrophys. J. Lett.*, **234**, L177–L181, (1979). 5.1.1, 8.1
- [142] Shapiro, S.L., and Teukolsky, S.A., “Gravitational Collapse to Neutron Stars and Black Holes: Computer Generation of Spherical Spacetimes”, *Astrophys. J.*, **235**, 199–215, (1980). Related online version (cited on 05 February 2007):
<http://adsabs.harvard.edu/abs/1980ApJ...235..199S>. 5.1.1, 2, 8.1
- [143] Shapiro, S.L., and Teukolsky, S.A., “Relativistic stellar dynamics on the computer. I. Motivation and Numerical Method”, *Astrophys. J.*, **298**, 34–57, (1985). 5.1.1, 8.1
- [144] Shapiro, S.L., and Teukolsky, S.A., “Relativistic stellar dynamics on the computer. II. Physical applications”, *Astrophys. J.*, **298**, 58–79, (1985). 5.1.1, 8.1
- [145] Shapiro, S.L., and Teukolsky, S.A., “Collision of relativistic clusters and the formation of black holes”, *Phys. Rev. D*, **45**, 2739–2750, (1992). 8.2
- [146] Shibata, M., “Apparent horizon finder for a special family of spacetimes in 3D numerical relativity”, *Phys. Rev. D*, **55**, 2002–2013, (1997). 7.4, 8.2, 8.4, 8.5.3, 8.5.4, 8.5.7
- [147] Shibata, M., and Uryū, K., “Apparent Horizon Finder for General Three-Dimensional Spaces”, *Phys. Rev. D*, **62**, 087501, (2000). 8.5.1, 8.5.3, 8.5.4, 8.5.7
- [148] Shoemaker, D.M., *Apparent Horizons in Binary Black Hole Spacetimes*, Ph.D. Thesis, (The University of Texas at Austin, Austin, U.S.A., 1999). 8.7, 8.7.1, 8.7.2, 8.7.3
- [149] Shoemaker, D.M., Huq, M.F., and Matzner, R.A., “Generic tracking of multiple apparent horizons with level flow”, *Phys. Rev. D*, **62**, 124005, (2000). Related online version (cited on 09 January 2006):
<http://arXiv.org/abs/gr-qc/0006042>. 8.7, 8.7.1, 8.7.2, 8.7.3
- [150] Stoer, J., and Bulirsch, R., *Introduction to Numerical Analysis*, (Springer, Berlin, Germany; New York, U.S.A., 1980). 8.5.4
- [151] Szilágyi, B., Pollney, D., Rezzolla, L., Thornburg, J., and Winicour, J., “An explicit harmonic code for black-hole evolution using excision”, (2007). URL (cited on 09 April 2007):
<http://arXiv.org/abs/gr-qc/0612150>. 17

- [152] Teukolsky, S.A., “On the Stability of the Iterated Crank–Nicholson Method in Numerical Relativity”, *Phys. Rev. D*, **61**, 087501, (2000). 61
- [153] Thornburg, J., “Finding apparent horizons in numerical relativity”, *Phys. Rev. D*, **54**, 4899–4918, (1996). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9508014>. 7.4, 8.5.3, 8.5.4, 40, 8.5.4, 8.5.4, 8.5.7, 49
- [154] Thornburg, J., “A 3+1 Computational Scheme for Dynamic Spherically Symmetric Black Hole Spacetimes – I: Initial Data”, *Phys. Rev. D*, **59**, 104007, (1999). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9801087>. 28, 8.1, 8.1
- [155] Thornburg, J., “A 3+1 Computational Scheme for Dynamic Spherically Symmetric Black Hole Spacetimes – II: Time Evolution”, (1999). URL (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/9906022>. 8.1, 9
- [156] Thornburg, J., “A fast apparent horizon finder for three-dimensional Cartesian grids in numerical relativity”, *Class. Quantum Grav.*, **21**, 743–766, (2004). Related online version (cited on 09 January 2006): <http://arXiv.org/abs/gr-qc/0306056>. 3.2, 7.4, 7.5, 8.5.1, 8.5.2, 8.5.4, 8.5.5, 8.5.7, 43, 10, 8.6.4, 52, 8.6.6, 8.7.1
- [157] Tod, K.P., “Looking for marginally trapped surfaces”, *Class. Quantum Grav.*, **8**, L115–L118, (1991). 8.7, 8.7, 8.7.4
- [158] Trottenberg, U., Oosterlee, C.W., and Schüller, A., *Multigrid*, (Academic Press, San Diego, U.S.A., 2001). 44, 62
- [159] Čadež, A., “Apparent Horizons in the Two-Black-Hole Problem”, *Ann. Phys. (N.Y.)*, **83**, 449–457, (1974). 8.2
- [160] Wald, R.M., *General Relativity*, (University of Chicago Press, Chicago, U.S.A., 1984). 1, 1, 19
- [161] Wald, R.M., and Iyer, V., “Trapped surfaces in the Schwarzschild geometry and cosmic censorship”, *Phys. Rev. D*, **44**, R3719–R3722, (1991). 20
- [162] Walker, P., *Horizons, Hyperbolic Systems, and Inner Boundary Conditions in Numerical Relativity*, Ph.D. Thesis, (University of Illinois at Urbana-Champaign, Urbana, U.S.A., 1998). 5.3, 5.3.2, 12, 13
- [163] York Jr, J.W., “Kinematics and Dynamics of General Relativity”, in Smarr, L.L., ed., *Sources of Gravitational Radiation*, Proceedings of the Battelle Seattle Workshop, July 24–August 4, 1978, 83–126, (Cambridge University Press, Cambridge, U.K.; New York, U.S.A., 1979). I, 1, 1
- [164] York Jr, J.W., “Initial Data for Collisions of Black Holes and Other Gravitational Miscellany”, in Evans, C.R., Finn, L.S., and Hobill, D.W., eds., *Frontiers in Numerical Relativity*, Proceedings of the International Workshop on Numerical Relativity, University of Illinois at Urbana-Champaign, U.S.A., May 9–13, 1988, 89–109, (Cambridge University Press, Cambridge, U.K.; New York, U.S.A., 1989). 23