Applications of Numerical Relativity Beyond Astrophysics

This dissertation is submitted for the degree of Doctor of Philosophy

Saran Tunyasuvunakool
Emmanuel College

Supervisor: Pau Figueras

January 2017

This dissertation is dedicated to my wife, Kathryn, whose love and support I have valued dearly…

…and to my parents, who have been nurturing my sense of curiosity for as long as I can remember.

Abstract

Title: Applications of Numerical Relativity Beyond Astrophysics
Author: Saran Tunyasuvunakool

Numerical relativity has proven to be a successful and robust tool for non-perturbative studies of gravitational phenomena in the highly dynamical and/or non-linear regime. Perhaps the most prominent achievement in the field is the breakthrough success in simulating the merger of binary black hole systems. Gravitational waveforms resulting from these simulations serve as precise theoretical predictions of general relativity, which can be tested against observational data, such as those recently made by the LIGO experiment.

This dissertation explores applications of numerical relativity which lie beyond the realm of astrophysics. One motivation for this comes from the AdS/CFT correspondence, which allows us to study strongly coupled quantum field theories by considering classical gravity with a negative cosmological constant. More concretely, we construct stationary asymptotically anti-de Sitter spacetimes by numerically solving the Einstein equations in a strongly elliptic form, subject to various boundary conditions corresponding to the physical setting of interest. Three applications of this technique are presented here. 1) A toroidal "black ring" in global AdS5, which provides a more complete phase diagram for AdS5 black holes. 2) A black hole on an AdS soliton background, which is dual to a localised ball of deconfined plasma surrounded by confined matter. 3) A rotating horizon extending to the AdS boundary, which allows us to the study the behaviour of the CFT in the presence of a rotating black hole.

Outside of AdS/CFT, time-dependent numerical relativity in higher dimensions can also inform inquiries into the mathematical properties of general relativity as a theory of gravity. In particular, long, thin black hole horizons are known to be subject to the Gregory–Laflamme instability, and this is expected to result in an eventual violation of the weak cosmic censorship conjecture. A landmark simulation of the black string confirmed this in the Kaluza–Klein setting, however the generalisation of this setup to asymptotically flat black rings poses new challenges for numerical relativity. Even after a successful simulation, the resulting apparent horizons possess nontrivial geometries which are problematic for existing horizon finding methods. This dissertation also covers aspects of technical development in the GRChombo adaptive mesh refinement code which were necessary for the successful evolution and analysis of a black ring instability.

Acknowledgements

This dissertation is—in large parts—the product of the experience, patience and guidance of my supervisor, Pau Figueras. Through him, I have learned to become a better physicist and, ultimately, a better learner. I am also indebted to Prof. Harvey Reall for accepting me as a PhD student in the first instance, and for introducing me to Pau as I expressed an interest in numerical work.

I am extremely grateful to my fellow PhD student Markus Kunesch, whose tireless effort and enthusiasm has been indispensable in riding out the rougher waves of our time-evolution work. I would also like to thank Katy Clough, Will Cook, Hal Finkel, Eugene Lim, and Ulrich Sperhake, for our thoroughly enjoyable and stimulating collaborations.

My research was almost as much software engineering as it was gravitational physics, and on this side I have benefited tremendously from the friendship and the numerous, lengthy, and insightful discussions with Juha Jäykkä, Kacper Kornet, and James Briggs. Without their hard work, the smooth operation of the COSMOS supercomputer would have been impossible over the past few years. I must also acknowledge the hospitality of Hal Finkel during my visit to the Argonne Leadership Computing Facility, which was one of the most memorable trips of my PhD tenure.

Throughout my time at DAMTP, I have greatly enjoyed the company of, among others, Chris Blair, Tom Eaves, Dejan Gajic, Davide Gerosa, Rahul Jha, Joe Keir, Emanuel Malek, Giuseppe Papallo, Alasdair Routh, Andrew Singleton, Carl Turner, and Helvi Witek. I am also very thankful for the longstanding friendship of Low Zhen Lin and Stephen Shaw.

My PhD was supported by the European Research Council (ERC), initially through grant No. ERC-2011-StG 279363-HiDGR and later through grant No. ERC-2014-StG 639022-NewNGR. The research presented in this dissertation was made possible through the computational resources provided by the COSMOS Shared Memory system at DAMTP, University of Cambridge, operated on behalf of the STFC DiRAC HPC Facility (funded by BIS National E-infrastructure capital Grant No. ST/J005673/1 and STFC Grants No. ST/H008586/1, No. ST/K00333X/1), the SuperMike-II cluster at the Louisiana State University, and the Stampede cluster at the Texas Advanced Computing Centre (TACC), through the NSF-XSEDE Grant No. PHY-090003.

Declarations

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared below and specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution.

Chapter 3 of this dissertation is based on the following co-authored publication:

P. Figueras and S. Tunyasuvunakool,
"CFTs in rotating black hole backgrounds," Classical and Quantum Gravity 30 (2013) 125015

Chapter 4 of this dissertation is based on the following co-authored publication:

P. Figueras and S. Tunyasuvunakool,
"Localized Plasma Balls," Journal of High Energy Physics 06 (2014) 025

Chapter 5 of this dissertation is based on the following co-authored publication:

P. Figueras and S. Tunyasuvunakool,
"Black rings in global anti-de Sitter space," Journal of High Energy Physics 03 (2015) 149

Chapters 6 and 8 of this dissertation contains results from the following co-authored publications:

K. Clough, P. Figueras, H. Finkel, M. Kunesch, E. A. Lim, and S. Tunyasuvunakool,
"GRChombo: Numerical Relativity with Adaptive Mesh Refinement,"
Classical and Quantum Gravity 32 no. 24, (2015) 245011

P. Figueras, M. Kunesch, and S. Tunyasuvunakool,
"End Point of Black Ring Instabilities and the Weak Cosmic Censorship Conjecture,"
Physical Review Letters 116 no. 7, (2016) 071102

Contents

Chapter 1

Preliminaries

1.1 Introduction

General relativity (GR) has been an extremely successful theory of gravitation. Although the geometric nature of the Einstein equation provides an elegant and simple description of gravity, its nonlinearity means that numerical methods are often required to investigate physical phenomena in the strong-field regime. One early result of numerical relativity (NR) was Choptuik's discovery of critical phenomena in the gravitational collapse of a massless scalar field [1]. The most noteworthy application of numerical relativity, however, is probably the simulation of black hole binaries. Initial progress in simulating black holes occurred as part of the *Binary Black Hole Grand Challenge* [2], and was spurred by the anticipated need for theoretical waveform predictions in a successful detection of gravitational waves at the then-upcoming Laser Interferometer Gravitational Wave Observatory (LIGO). A full 'inspiral-merger-ringdown' simulation of a binary black hole system only became possible after over a decade of theoretical work on formulations of the Einstein equation, gauge conditions, and initial conditions, along with advances in simulation code and computing power. Remarkably, the eventual success in long-term stable evolution of such a system was achieved through two distinct numerical schemes by three independent groups [3, 4, 5]. Indeed, when gravitational waves were finally detected earlier this year at LIGO [6, 7], the events were found to be consistent with waveforms generated via numerical simulations of binary black hole mergers.

The maturation of numerical techniques in general relativity has led to their wider application in

other areas of gravitational research, beyond the usual astrophysical setting. One important area is the study of the mathematical properties of GR itself. In the non-dynamical regime, black holes are the most fundamental stationary objects in the theory. Through the celebrated uniqueness theorems [8, 9, 10, 11], we know that asymptotically flat black holes in 4 dimensions are described by the Kerr metric [12]. In higher dimensions, the Kerr solution generalises into the Myers–Perry black holes [13], however they are no longer unique. The discovery of the five-dimensional black ring [14] shows that there can be multiple black hole solutions of different horizon topologies at a given mass and angular momentum. Following this work, many other 'black objects' have been found analytically in five dimensions [15, 16, 17, 18, 19, 20]. However, in six or more dimensions, even the black rings are only known numerically [21, 22]. While constraints exist on a black hole horizon's topology [23], the complete classification of black holes in higher dimensions remains an open problem. Another major motivation to construct novel stationary solutions comes from the AdS/CFT correspondence [24], which allows us to study properties of certain strongly coupled quantum conformal field theories (CFT) by constructing asymptotically anti-de Sitter (AdS) solutions to classical gravity, subject to appropriate boundary conditions. Analytic methods have successfully been used to construct new solutions, but they are subject to limitations. Exact solution generating techniques (e.g. inverse scattering transform [25, 26], algebraically special solutions [27]) are limited to geometries with a high degree of symmetry, while approximate techniques (e.g. blackfolds [28], large D expansion [29]) are limited to certain corners of the parameter space. Numerical methods therefore provide an invaluable complementary framework in which to investigate stationary solutions which remain inaccessible via traditional analytic methods.

The programme to construct new black hole solutions goes hand in hand with the investigation of their dynamical stability. In astrophysics, the role of the Kerr black holes as the generic end state of gravitational collapses is dependent on them being nonlinearly stable. While a rigorous mathematical proof of this remains a major open problem in mathematical relativity, there is good evidence that Kerr black holes should indeed be stable. In higher dimensions, however, the situation is less clear, especially since the pioneering work of Gregory and Laflamme [30] showed that black strings and black branes are linearly unstable to gravitational perturbation. In this context, NR has proven to be a useful tool for studying the nonlinear evolution of perturbed, higher dimensional black holes. The landmark simulation of the black string [31] revealed the self-similar dynamics

of the Gregory–Laflamme (GL) instability. The formation of a fractal-like structure on the black string's horizon provided strong evidence for a violation of the weak cosmic censorship conjecture [32] in asymptotically Kaluza–Klein spacetimes. In the asymptotically flat setting, simulations of the Myers–Perry black hole in $D = 6$ uncovered a new type of instability which causes the black hole to lose angular momentum and settle into a state of slower rotation [33]. Furthermore, thin black rings and ultraspinning Myers–Perry black holes are also expected to be subject to the GL instability. The application of numerical relativity to these systems would allow us to study the nonlinear interaction between these instabilities and follow them to the end state.

In this dissertation, we explore a number of different applications of numerical methods to the study of general relativity, beyond the realm of astrophysical simulations. The dissertation is divided into two main parts. The first part, consisting of Chapters 2,3,4 and 5, concerns the construction of various stationary black hole solutions in AdS. Through AdS/CFT, the solutions that we construct are dual to a CFT in the background of rotating black holes, a localised ball of deconfined plasma surrounded by a confining vacuum, and a new CFT equilibrium phase. The second part, consisting of Chapters 6,7 and 8, concerns the dynamical simulation of higher-dimensional spacetimes. Through simulations of black rings, we found a new type of instability, and evidence for a possible violation of the weak cosmic censorship conjecture in the asymptotically flat setting in $D = 5$. While these two problem areas pose different challenges from a practical viewpoint, they both require the Einstein equation to be cast in a form which is suitable for numerical computations. The methods that we use throughout this dissertation are based on the same basic covariant modification of the Einstein equation, which removes the pure diffeomorphism terms from the principle part of the PDE. We devote the rest of this introductory chapter to providing a high-level overview of this general idea, while deferring more specific remarks and details to later chapters.

1.2 Gauge freedom in the Einstein equation

General relativity is most naturally phrased in its covariant, tensorial form. In D dimensions, in the presence of a cosmological constant Λ and matter with energy-momentum tensor T_{ab}, spacetime

is curved according to the Einstein field equation

$$R_{ab} - \frac{2}{D-2}\Lambda g_{ab} = 8\pi\left(T_{ab} - \frac{1}{D-2}T^{c}{}_{c}\,g_{ab}\right), \tag{1.1}$$

where g_{ab} is the metric on the spacetime manifold $\mathcal{M}$, and R_{ab} is its Ricci tensor. Given a coordinate chart $x^{\mu} = \left(x^{0}, \ldots, x^{D-1}\right)$ on $\mathcal{M}$, the components of the Ricci tensor are given by

$$R_{\mu\nu} = \partial_{\rho}\Gamma^{\rho}_{\mu\nu} - \partial_{\mu}\Gamma^{\rho}_{\rho\nu} + \Gamma^{\rho}_{\rho\sigma}\Gamma^{\sigma}_{\mu\nu} - \Gamma^{\rho}_{\sigma\mu}\Gamma^{\sigma}_{\rho\nu}, \tag{1.2}$$

$$\Gamma^{\mu}_{\nu\rho} = \frac{1}{2}g^{\mu\sigma}\left(\partial_{\rho}g_{\nu\sigma} + \partial_{\nu}g_{\rho\sigma} - \partial_{\sigma}g_{\nu\rho}\right). \tag{1.3}$$

Since both g_{ab} and R_{ab} are symmetric tensors, (1.2) defines a system of $\frac{1}{2}D\,(D+1)$ nonlinear PDEs for $\frac{1}{2}D\,(D+1)$ metric components $g_{\mu\nu}$, which can be regarded as unknown functions to be solved for. However, the Ricci tensor satisfies the *contracted Bianchi identity*, which serves to constrain its components via

$$\nabla_{b}R^{b}{}_{a} - \frac{1}{2}\nabla_{a}R = 0. \tag{1.4}$$

The presence of D constraints implies that the Ricci tensor possesses D fewer independent components than the metric, and the PDE system is therefore underdetermined. This should not be surprising given the coordinate invariance arising from the geometric nature of Einstein's equation. Given a spacetime $(\mathcal{M}, g_{ab})$, there are uncountably many systems of coordinates by which it can be covered. Therefore, given a set of metric components $g_{\mu\nu}$, one can perform an arbitrary coordinate transformation $\hat{x}^{\alpha} := \hat{x}^{\alpha}(x^{\mu})$ to obtain a different set of metric components $\hat{g}_{\alpha\beta}$ describing a physically identical spacetime. In other words, while (1.1) is sufficient to determine the *geometry*, it does *not* uniquely determine the values of the metric components. To fix these, we must additionally enforce some conditions on the coordinates. The D additional coordinate conditions would then supplement the D degrees of freedom that were rendered redundant through (1.4).

More precisely, the tensorial form of the Einstein equation, as a system of PDE, is not manifestly strongly hyperbolic. Therefore, the resulting initial value problem is not necessarily well-posed. We can analyse the character of the PDE system (1.2) by first recasting the Ricci tensor components

(1.2) into the so-called de Donder form [34]

$$\begin{aligned} R_{\mu\nu} = &-\frac{1}{2}g^{\rho\sigma}\partial_\rho\partial_\sigma g_{\mu\nu} + \partial_{(\mu}\Gamma_{\nu)} \\ &+ g_{\rho\sigma}g^{\tau\lambda}\Gamma^\rho_{\tau\mu}\Gamma^\sigma_{\lambda\nu} - \Gamma_\rho\Gamma^\rho_{\mu\nu} - \Gamma^\rho_{\sigma\tau}g_{\rho(\mu}\partial_{\nu)}g^{\sigma\tau} \end{aligned} \tag{1.5}$$

where we define $\Gamma^\mu := g^{\nu\rho}\Gamma^\mu_{\nu\rho}$ and $\Gamma_\mu := g_{\mu\nu}\Gamma^\nu$. The principal part of $R_{\mu\nu}$ is contained in the first two terms of (1.5). The first term can be recognised as the standard wave operator on the Lorentzian manifold, which is manifestly symmetric hyperbolic. It is therefore the second term that is problematic for the hyperbolicity of the system. By treating each coordinate x^μ as a function on the spacetime, we obtain

$$\Box x^\mu = g^{\nu\rho}\nabla_\nu\nabla_\rho x^\mu = g^{\nu\rho}\nabla_\nu\delta^\mu_\rho = -\Gamma^\mu. \tag{1.6}$$

Therefore, the term $\partial_{(\mu}\Gamma_{\nu)}$ in the principal part of the Einstein equation is only associated with properties of the coordinate system and not the geometry of spacetime.

This dissertation concerns the construction of both stationary and dynamical solutions to the Einstein equation. In both cases, the numerical methods used follow the same basic principle to fix the 'pure gauge' degrees of freedom by modifying the Einstein equation itself. More precisely, we consider solving the following equation instead of (1.1):

$$R_{ab} + \nabla_a Z_b + \nabla_b Z_a - \frac{2}{D-2}\Lambda g_{ab} = 8\pi\left(T_{ab} - \frac{1}{D-2}T^c{}_c g_{ab}\right). \tag{1.7}$$

The idea is to choose the vector Z such that the subspace of solutions to (1.7) with $Z \equiv 0$ coincide exactly with the space of solutions to (1.1). In the context of stationary solutions, this leads to the *Einstein–DeTurck method*, as reviewed in Section 2.2. The dynamical analogue of this is the *generalised harmonic coordinates* (GHC) evolution scheme, which was successfully used in the first breakthrough binary black hole simulation [3]. However, current GHC codes require black holes to be excised from the computational domain, which is a difficult procedure to implement in practice. An alternative is perform a '$d+1$ decomposition' (where d denotes the number of spatial dimensions) of (1.7) and use the *moving puncture gauge*. A conformally-decomposed version of this, known as CCZ4, is reviewed in Section 6.2.

CHAPTER 2

Numerical Construction of Stationary Black Holes in AdS

2.1 The AdS/CFT correspondence

The AdS/CFT correspondence is one of the most influential developments in theoretical physics in recent decades. Formulated by Maldacena in 1997 [24], AdS/CFT is a conjectured duality between certain conformal field theories (CFTs) in D dimensions and string theory—a theory of quantum gravity—on $D+1$-dimensional asymptotically anti-de Sitter (AdS) spacetimes. While the complete equivalence of these two theories still holds the status of a conjecture, research over the years has shown that the correspondence holds up to many nontrivial tests. From a purely theoretical perspective, given that a precise definition of quantum gravity itself is an open problem, one intriguing possibility arising from this duality is that it would allow us to define quantum gravity in terms of a gauge theory, which is far better understood. From a more practical viewpoint, since classical general relativity corresponds to the strong coupling regime on the CFT side, the duality allows us to probe the behaviour of quantum field theories where perturbative techniques do not apply. In this section, we only provide a brief, high-level overview of the topic, so as to motivate our interest in constructing stationary black hole solutions in AdS numerically. Many excellent review articles exist which provide a more detailed theoretical treatment of the subject.

The original derivation of AdS/CFT [24] considers a stack of N coincident D3-branes in 10-dimensional type IIB string theory. In the weak coupling regime $g_s N \ll 1$, where g_s is the string coupling constant, the spacetime curvature induced by the D-branes is negligible. In the low-energy

limit $l_s \to 0$, where l_s is the string length, the string excitation modes on the branes decouple from those in the bulk. The effective action on the branes is that of the $\mathcal{N} = 4$ super Yang–Mills (SYM) conformal field theory with an $SU(N)$ gauge group in 4D flat space, while the bulk action is that of supergravity in 10D flat space. On the other hand, at strong coupling $g_s N \gg 1$, the D-branes backreact strongly on the spacetime. In this regime, the system is described by the extremal black 3-brane solution of supergravity [35]

$$\begin{gathered} \mathrm{d}s^2 = f(r)^{-1/2}\,\eta_{ij}\,\mathrm{d}x^i\,\mathrm{d}x^j + f(r)^{1/2}\left(\mathrm{d}r^2 + r^2\mathrm{d}\Omega_5^2\right) \\ f(r) \coloneqq 1 + \frac{4\,\pi\,g_s\,N\,l_s^4}{r^4} \end{gathered}\quad , \tag{2.1}$$

where the μ, ν indices range over the 4 coordinates on the brane, and $\mathrm{d}\Omega_5^2$ is the standard round metric on S^5. Like any other extremal black hole solution, the horizon at $r = 0$ lies at the end of an infinite 'throat'. Defining the characteristic length scale $\ell \coloneqq (4\,\pi\,g_s\,N)^{1/4}\,l_s$, we can consider the *near horizon geometry* of the black brane, in the limit of $r \ll \ell$. This is given by the metric

$$\mathrm{d}s^2 = \frac{r^2}{\ell^2}\,\eta_{ij}\,\mathrm{d}x^i\,\mathrm{d}x^j + \frac{\ell^2}{r^2}\,\mathrm{d}r^2 + \ell^2\,\mathrm{d}\Omega_5^2. \tag{2.2}$$

Defining a new coordinate $z \coloneqq \ell^2/r$, this can be recognised as the metric on the Poincaré patch of $\mathrm{AdS}_5 \times S^5$

$$\mathrm{d}s^2 = \frac{\ell^2}{z^2}\left(\eta_{ij}\,\mathrm{d}x^i\,\mathrm{d}x^j + \mathrm{d}z^2\right) + \ell^2\,\mathrm{d}\Omega_5^2. \tag{2.3}$$

The length scale ℓ is usually referred to as the *AdS radius*. In the low-energy limit $l_s \to 0$ (equivalently, $\alpha' \coloneqq l_s^2 \to 0$), the excitation modes approaching the horizon decouple from those propagating in the asymptotic region away from the throat. The former sector is described by supergravity on $\mathrm{AdS}_5 \times S_5$, while the latter is again described by supergravity in 10D flat space.

Although these two low-energy descriptions of the D3-branes system apply to two distinct regimes of $g_s N$, in principle the $\mathcal{N} = 4$ SYM gauge theory is well-defined at any coupling. This leads to the remarkable conclusion that the $\mathcal{N} = 4$ $SU(N)$ SYM, which is a non-gravitational system, is equivalent to supergravity on $\mathrm{AdS}_5 \times S^5$. Maldacena went further to conjecture that this correspondence extends beyond the supergravity limit into one between $\mathcal{N} = 4$ SYM and the full IIB string theory on $\mathrm{AdS}_5 \times S^5$. Amongst the arguments in support of this conjecture, the most immediate is perhaps the fact that the two theories share the same symmetries. The isometry group

of AdS_5 is $SO(4,2)$ while the isometry group of S^5 is $SO(6)$. From the CFT perspective, the former appears as the conformal group in four dimensions, while the latter appears as the group of R-symmetry between the scalars and the fermions. The radial coordinate r in (2.2) can be identified with the energy scale E of the CFT. Specifically, the conformal boundary $z = 0$ corresponds to the high energy, short wavelength 'UV' end. The CFT may be regarded as living on the transverse directions x^μ, whose geometry belongs to the same conformal class as the AdS boundary. It is therefore common parlance to refer to the CFT as *living on the boundary of AdS*.

The correspondence also provides a relationship between the parameters of the two theories involved. On the 'AdS' side of the correspondence, we have the string coupling g_s, the string length l_s, and the AdS radius ℓ, while on the 'CFT' side we have the Yang–Mills coupling g_{YM} and the rank N of the gauge group $SU(N)$. Through AdS/CFT, these parameters are related via

$$\begin{aligned} 4\,\pi\, g_s &= g_{\mathrm{YM}}^2 \\ \frac{\ell^4}{l_s^4} &= 4\,\pi\, g_s\, N = g_{\mathrm{YM}}^2\, N \coloneqq \lambda, \end{aligned} \tag{2.4}$$

where $\lambda \coloneqq g_{\mathrm{YM}}^2\, N$ is the *t'Hooft coupling*. The $SU(N)$ gauge theory has a smooth t'Hooft limit [36] $N \to \infty$ at fixed λ, with a perturbative expansion in $1/N$. On the AdS side, this corresponds to taking the $g_s \to 0$ limit, which gives us classical gravity with a perturbative expansion in g_s. The correspondence therefore allows us to study the behaviour of a strongly coupled quantum CFT, at least to leading order in $1/N$, by performing calculations in classical gravity. It should also be noted that the duality between $\mathcal{N} = 4$ SYM in $D = 4$ and supergravity in $AdS_5 \times S^5$ is only one example of the AdS/CFT correspondence. More generally, the correspondence applies in any number of dimensions D, and it is in this more general context that we continue our discussion.

In order to utilise this correspondence to perform calculations, we must be able to relate CFT observables to quantities in the AdS bulk: the so-called *AdS/CFT dictionary*. This was first proposed in [37, 38], where the value of the supergravity action on AdS was equated with the generating functional of CFT operators. The results relate IR-divergent quantities on the supergravity side with UV-divergent quantities on the CFT side. Noting that this should be regarded as a relationship between bare quantities, the authors of [39] provided a systematic means to renormalise the theories in order to obtain a relationship between finite quantities on both sides; a prescription

usually referred to as *holographic renormalisation*. Of particular interest to us is their formula for the expectation value of the stress-energy tensor of a CFT that admits a gravitational dual. In the neighbourhood of the conformal boundary $z = 0$, any given asymptotically AdS_{D+1} spacetime may be expressed in *Fefferman–Graham coordinates*

$$\mathrm{d}s^2 = \frac{\ell^2}{z^2}\left(\mathrm{d}z^2 + g_{ij}(x,z)\,\mathrm{d}x^i\mathrm{d}x^j\right), \tag{2.5}$$

where crucially $g_{\mu\nu}$ admits an expansion in terms of z in the form

$$g(x,z) = g_{(0)}(x) + g_{(2)}(x)\,z^2 + \ldots + g_{(D)}(x)\,z^D + \mathbb{1}_{D\equiv 0\,(\mathrm{mod}2)}\,h_{(D)}(x)\,z^D\,\log z^2 + \mathcal{O}(z^{D+1}), \tag{2.6}$$

where the logarithmic term only exists when D is even. By inserting the metric (2.5) into the Einstein equation (1.1) and solving it at each order in z, we can immediately show that all the $g_{(n)}$ are fully determined for $n < D$. However, the Einstein equation only fixes the divergence and trace of $g_{(D)}$, and the dual CFT data is required to determine it fully. In terms of the FG expansion, the expectation value of the dual CFT stress-energy tensor is given by

$$\langle T_{ij}\rangle = \frac{D\,\ell^{D-1}}{16\,\pi\,G_{\mathrm{N}}}\,g_{(D)ij} + X_{ij}[g_{(D)}], \tag{2.7}$$

where $G_{\mathrm{N}} = g_s^2\,l_s^{D-2}$ is Newton's constant, and X_{ij} is an expression which depends on the number of dimensions D. For odd D, we have $X_{ij} = 0$ and $\mathrm{tr}\,g_{(D)} = 0$, consistent with the absence of conformal anomalies. For even D, [40] provides expressions for $D = 2, 4, 6$.

In chapters 3, 4 and 5, we apply AdS/CFT to investigate certain properties of CFTs in the strongly coupled regime. Through AdS/CFT, this is reduced to the problem of constructing stationary solutions to the Einstein equation, subject to boundary conditions defined by the CFT physics of interest. Chapter 3 considers the behaviour of CFTs when placed in the background of a rotating black hole, which is relevant to the study of Hawking radiation of strongly coupled field theories. Chapter 4 builds upon the programme of [41] to study the deconfinement phase transition through AdS/CFT, by constructing a solution which is dual to a localised ball of deconfined CFT plasma surrounded by a confining vacuum. This can be viewed as a natural extension of the domain wall solution previously constructed in [42]. Chapter 5 concerns the construction of an AdS analogue to the well-known Emparan–Reall black ring solution [14]. Since these can be taken to have the

same temperature as the rotating AdS black hole solutions [43], they can be regarded as the dual of a new CFT phase.

The remainder of this chapter is devoted to a formulation of the Einstein equation suitable for numerically constructing stationary spacetimes, as well as to the methods used to solve the resulting PDE system.

2.2 The Einstein–DeTurck method

In many cases, stationary solutions to the Einstein equation can be regarded as the end point of some dynamical gravitational process. In principle, then, one could simply construct such solutions by preparing some appropriate initial configuration and evolving it until a steady state is reached. Clearly, this will not allow us to control the final solution directly. Instead, we must find an initial condition which dynamically evolves into the desired solution. Generally speaking, this is in itself a non-trivial problem. Another major limitation is that many stationary solutions of interest are dynamically unstable, and therefore we will not be able to reach these solutions through time evolution. The black ring is a good example of such an unstable solution.

In this dissertation, we use the Einstein–DeTurck method which recasts the Einstein equation into a strongly elliptic form, thus allowing us to construct stationary solutions directly by solving a boundary value problem. This chapter provides an overview of this method, and a much more extensive review can be found in [44]. We begin by specifying some fixed *background metric* $\bar{g}_{ab}$ on the spacetime, with connection components $\bar{\Gamma}^{\mu}_{\nu\rho}$. This will amount to specifying a particular choice of coordinates for our solution. We then define the *Ricci-DeTurck tensor*

$$\begin{gathered} R^{H}_{ab} \coloneqq R_{ab} - \nabla_{(a}\xi_{b)}, \\ \text{and the components of } \xi^{a} \text{ are given by } \xi^{\mu} \coloneqq g^{\nu\rho}\left(\Gamma^{\mu}_{\nu\rho} - \bar{\Gamma}^{\mu}_{\nu\rho}\right). \end{gathered} \tag{2.8}$$

The superscript H here stands for *harmonic*. Note that ξ^{a} is a *bona fide* vector field, as it is a difference of two connections. Replacing R_{ab} in the Einstein equation (1.1) with R^{H}_{ab} yields the

harmonic Einstein equation

$$R_{ab} - \nabla_{(a}\xi_{b)} - \frac{2}{D-2}\Lambda g_{ab} = 8\pi\left(T_{ab} - \frac{1}{D-2}T^{c}{}_{c}\,g_{ab}\right). \tag{2.9}$$

Solutions to (2.9) also satisfy (1.1) provided that $\xi^a \equiv 0$. This in turn imposes the *generalised harmonic condition* $\Box x^\mu = -\bar{\Gamma}^\mu$ on the coordinates. This modification directly subtracts the $\partial_{(\mu}\Gamma_{\nu)}$ term off (1.5), making the principal part of the harmonic Einstein equation identical to that of the Laplacian. We will now argue that the harmonic Einstein operator can be consistently restricted to the space of stationary metrics, and that by doing so, it becomes an elliptic operator which is suitable for the numerical construction of stationary spacetimes.

The DeTurck method was first proposed in [45] as a robust way to solve static problems. In that case, the metric trivially extends to a Riemannian signature, and the wave operator in the principal part of R^H_{ab} is manifestly elliptic. It was later shown in [46] that the strong ellipticity of the harmonic Einstein equation applies beyond the static regime to a much larger class of stationary solutions. We will only present the argument in its most general form here, which is valid for stationary black holes whose Killing vector field of stationarity is not necessarily globally timelike. These include, for example, rotating black hole solutions with an ergoregion. The argument is motivated by the rigidity property of stationary black holes [47]. Let $(\mathcal{M}, g_{ab})$ be a stationary analytic, asymptotically flat/AdS vacuum black hole spacetime. Let T be the Killing vector field of stationarity, i.e. one which is asymptotically timelike, then T can either be a horizon generator or not. In the latter case, the rigidity theorem states that there exist at least one Killing vector fields $R_{(\alpha)}$ ($\alpha = 1, \ldots, N$), each having a closed orbit of period 2π, such that $\{T, R_{(\alpha)}\}$ are mutually commuting. Furthermore, there exist some constants $\Omega^{(\alpha)}$ such that the Killing vector field $K := T + \sum_\alpha \Omega^{(\alpha)} R_{(\alpha)}$ is a horizon generator. Physically, this implies that the black hole rotates rigidly with angular velocities $\Omega^{(\alpha)}$ on each rotational plane, with respect to the static observer in the asymptotic region. We can then define an isometry-adapted coordinate system $y^A := \{t, y^\alpha\}$ such that $T = \partial/\partial t$ and $R_{(\alpha)} = \partial/\partial y^\alpha$. In these coordinates, the general stationary metric takes the form

$$\mathrm{d}s^2 = G_{AB}(x)\left(\mathrm{d}y^A + A^A_i(x)\,\mathrm{d}x^i\right)\left(\mathrm{d}y^B + A^B_j(x)\,\mathrm{d}x^j\right) + h_{ij}(x)\,\mathrm{d}x^i\mathrm{d}x^j, \tag{2.10}$$

where we use uppercase Latin indices A, B to range over isometry-adapted coordinates, and lower-

case Latin indices i, j to range over the remaining spatial directions. To make progress, we make an additional *assumption* that h_{ij} are components of a smooth metric with Riemannian signature, i.e. $\det h_{ij} \geq 0$, eveywhere in the exterior of the horizon. Since $\det g_{ab} = \det G_{AB} \det h_{ij}$, this implies that $\det G_{AB} \leq 0$, and therefore that there exists some linear combination of $\{T, R_{(\alpha)}\}$ which is timelike. Inserting this into the Ricci-DeTurck tensor, we find

$$R_{AB}^{H\,(\text{principal})} \simeq -\frac{1}{2} h^{kl}\, \partial_k \partial_l G_{AB}(x) \tag{2.11}$$

$$R_{Ai}^{H\,(\text{principal})} \simeq -\frac{1}{2} h^{kl}\, \partial_k \partial_l \left(G_{AB}(x)\, A_i^B(x) \right) \tag{2.12}$$

$$R_{ij}^{H\,(\text{principal})} \simeq -\frac{1}{2} h^{kl}\, \partial_k \partial_l \left(h_{ij}(x) + G_{AB}(x)\, A_i^A(x)\, A_j^B(x) \right). \tag{2.13}$$

Most importantly, the principal part of the harmonic Einstein equation is now the Laplacian on a Riemannian manifold with metric h_{ij}, and the equation is therefore elliptic.

Since we will be solving the nonlinear PDE system by iterative numerical methods, we must also require R_{ab}^H to share the same symmetries as the metric g_{ab}. This ensures that the metric remains within the family (2.10) at the end of each iteration. To achieve this, we need only require that the reference metric $\bar{g}_{ab}$ also possesses $\{T, R_{(\alpha)}\}$ as Killing vectors and satisfies the same assumption as g_{ab}. In practice, this means that the background metric components also belong to the family (2.10).

It is important to note, however, that Einstein metrics form a only subset of solutions to (2.9). Generally speaking, it is entirely possible to construct a solution to (2.9) with $\xi \neq 0$. These are called *Ricci solitons*. The authors [48] applied a maximum principle to show that Ricci solitons cannot exist when solving for *static* solutions. The absence of Ricci solitons has very recently been generalised to a more general class of stationary spacetimes satisfying a '$t - \phi$' reflection symmetry [49]. From a more practical point of view, however, we can verify *a posteriori* that our solution does indeed have vanishing ξ. Although the components of ξ will never be exactly zero in numerical solutions, the elliptic character of (2.9) nevertheless ensures that any solution obtained is locally unique, and so a Ricci soliton solution cannot occur arbitrarily close to an Einstein solution. We can therefore be reasonably certain that a numerical solution is *not* a Ricci soliton if the components of the vector ξ are both suitably small and exhibit convergence to zero close to machine precision limit as the spatial resolution is increased, at a rate appropriate to the

discretisation scheme employed.

In addition to a strongly elliptic system of PDEs, we also require an appropriate set of boundary conditions in order to form a well-posed problem. The boundaries of the computation domain can be either physical, such as the conformal boundary of AdS, or fictitious, such as those arising at the axes of rotational symmetries. Since the exact form of the boundary conditions depends on the metric ansatz used, we defer the discussion of boundary conditions for each specific problem to each of the subsequent chapters.

Finally, we remark that the DeTurck method has also been successfully used to construct stationary solutions which lie outside the class for which the aforementioned proof of strong ellipticity apply. In particular, the method seems to work when the spacetime only possesses a single, helical Killing field [50, 51], and even for 'flowing' geometries with a non-Killing horizon [52, 53] which are useful for studying heat transport in AdS/CFT. In spite of the lack of a theoretical guarantee of well-posedness, numerical convergence of the DeTurck method in these situations suggest that it is a very robust scheme for constructing stationary geometries indeed.

2.3 Numerical methods

Having formulated a boundary value problem, we now turn to the discussion of the methods which allow us to solve a nonlinear elliptic system numerically. To do this, the continuous problem domain must first be discretised on to a computational grid of finite spatial resolution. We therefore begin this section by discussing two discretisation schemes that are commonly used in the field, namely finite differences and pseudospectral collocation. The discretised problem is a nonlinear algebraic one, and we proceed to discuss two methods for solving this. The first of which is the Ricci flow method, whereby the elliptic problem is regarded as the steady state of an associated parabolic (diffusion) problem, and an initial guess is integrated in 'time' until a fixed point is reached. The second method is a line search based on Newton's method, whereby an the solution is iteratively corrected based on the linearised operator. This is the generalisation of the familiar Newton–Raphson root-finding method for 1D problems. Both of these methods, and indeed both discretisation schemes discussed, are used in the work presented in subsequent chapters of this

dissertation.

2.3.1 Spatial discretisation

Here, we will discuss spatial discretisation in one dimension, however the methods presented trivially generalise to arbitrarily many dimensions by simply performing the same procedure on each dimension. The most straightforward discretisation scheme is to construct the solution on grid points which are uniformly spaced at δx. In other words, on the interval $[x_L, x_R]$, our functions are evaluated at the points

$$x_i = x_L + i\,\frac{x_R - x_L}{N} \quad ; \quad i = 0, \ldots, N, \tag{2.14}$$

where the grid spacing is given by $\delta x := (x_R - x_L)/N$. We then replace continuous differentiation operators by *finite differences*: for a C^n-function, we evaluate derivatives at any particular grid point by fitting an nth degree polynomial onto the function values at the $n+1$ grid points closest to it. The value of the derivative is taken to be that of this interpolating polynomial. Here, n governs the rate of convergence in the continuum limit. More precisely, if we use an nth degree polynomial to evaluate the kth derivative, then we can expect the local error between our numerical solution and the true solution to decay like $\mathcal{O}(\delta x^{n-k+1})$ as $\delta x \to 0$. In practice, we typically specify a desired rate of convergence, and therefore would use a higher degree polynomial to evaluate a higher derivative accordingly.

The value of the approximated derivatives can be expressed as a weighted linear combination of function values at the $n+1$ closest grid points. The array of weights is referred to as the *stencil*, and we can regard the derivative as the discrete convolution of the function values with the stencil. In most of the computational domain, the resulting stencil will be symmetric around the point of interest. However, near domain boundaries, the stencil necessarily becomes asymmetric. This results in larger errors in these regions, as the stencil extends further away from the given point. Nevertheless, we should still expect errors to decay at the same rate everywhere as $\delta x \to 0$.

Indeed, we can contemplate fitting a single, high-degree polynomial onto the entire grid. However, this procedure is highly susceptible to spurious oscillations. This is the well-known *Runge's*

phenomenon, and becomes more severe as the degree of the interpolating polynomial increases. Nevertheless, it is possible to show that these spurious oscillations are minimised if we replace our uniformly spaced grid by a specific sets of points, known as *Chebyshev nodes*, canonically given on the interval $[-1, 1]$ by

$$x_i = -\cos\left(\frac{i\,\pi}{N}\right) \quad ; \quad i = 0, \ldots, N. \tag{2.15}$$

We are free to move and rescale these nodes to an arbitrary interval $[x_L, x_R]$ via

$$x_i = \frac{x_R + x_L}{2} - \frac{x_R - x_L}{2}\,\cos\left(\frac{i\,\pi}{N}\right) \quad ; \quad i = 0, \ldots, N. \tag{2.16}$$

The resulting grid has a higher density of points towards the two ends of the domain. On $[-1, 1]$, the Chebyshev nodes arise as roots of the *Chebyshev polynomials*, which are a basis of orthogonal polynomials defined via

$$T_N(\cos\theta) = \cos N\theta. \tag{2.17}$$

In the *pseudospectral collocation* scheme, we evaluate our functions on a grid of Chebyshev points, and fit the unique polynomial of degree N on to the entire dataset. Any polynomial $P_N(x)$ of degree N can be uniquely expressed as a sum of lower degree Chebyshev polynomials

$$P_N(x) = \sum_{k=0}^{N} c_k\, T_k(x). \tag{2.18}$$

On the other hand, since the Chebyshev polynomials form a complete basis, any function $f(x)$ can be expanded as an infinite Chebyshev series

$$f(x) = \sum_{k=0}^{\infty} c_k\, T_k(x). \tag{2.19}$$

We may therefore regard our discretisation scheme as a truncation of (2.19) to the first N terms. Through (2.17), we can see that the expansion (2.19) in x is related to a Fourier expansion in $\theta = \cos^{-1} x$. It is a standard result in Fourier analysis that if $f(x)$ is smooth, in the sense that it has infinitely many continuous derivatives, then the coefficients c_k decay faster than $\mathcal{O}(|k|^{-m})$ for any integer m, as $k \to \infty$. Therefore, the truncation error decays at a faster-than-polynomial rate as $N \to \infty$. This is usually referred to as *exponential convergence*, however actual exponential

rate of convergence is only guaranteed when the function is *analytic*. On the other hand, if $f(x)$ only has $p-1$ continuous derivatives, then we can only expect the truncation error in P_N to decay at $\mathcal{O}(N^{-p-1})$. In many cases, the pseudospectral method can be expected to converge much more rapidly than finite differences, and therefore it is usually possible to achieve acceptable accuracy using fewer grid points.

In order to evaluate derivatives, we simply differentiate the RHS of (2.18). By orthogonality, the coefficients c_k can be expressed as a linear combination of the function values on all the Chebyshev nodes. The resulting differentiation stencil therefore extends to cover the entire computational domain. In both the finite difference and pseudospectral schemes, the differentiation of the discretised function can be expressed as a matrix multiplication,

$$f'(x_i) = \sum_{j=0}^{N} D_{ij} f(x_j), \tag{2.20}$$

where D_{ij} is called the *differentiation matrix*. In the finite differences method, D_{ij} is a sparse band matrix, with the number of nonzero elements growing as $\mathcal{O}(N)$ in the number of grid points. In the pseudospectral method, D_{ij} is a dense matrix, with the number of nonzero elements growing as $\mathcal{O}(N^2)$. Therefore, while the faster convergence of the pseudospectral method allows us to use fewer grid points, the extra resources required by the dense differentiation matrix can sometimes offset its potential advantage. This is especially true in a distributed-memory parallel code, where the density of the differentiation matrix significantly increases communication cost and impedes scalability.

2.3.2 Numerical solutions to nonlinear elliptic systems

Spatial discretisation turns our system of nonlinear PDEs into a system of nonlinear algebraic equations, and in this subsection we discuss the numerical methods which can be used to solve them. In our first method, we begin by considering a family of metrics $g_{\mu\nu}$, parametrised by a single continuous parameter λ, with associated Ricci tensors $R_{\mu\nu}$. These metrics form a *Ricci flow* if they satisfy

$$\frac{\partial}{\partial\lambda} g_{ab}(\lambda; x) = -2R_{ab}(\lambda; x). \tag{2.21}$$

Ricci flow can be viewed as a geometric diffusion equation, where local curvature as described by the Ricci tensor is smoothed out over the *Ricci flow time* λ. However, the weak ellipticity of R_{ab} translated to weak parabolicity of the Ricci flow equation. In [54], DeTurck proved that by adding a term of the form (2.8), the flow can be rendered into the strongly parabolic *Ricci–DeTurck flow*

$$\frac{\partial}{\partial\lambda}\, g_{ab}(\lambda;x) = -2R_{ab}(\lambda;x) - 2\,\nabla_{(a}\xi_{b)}(\lambda;x) \coloneqq -2R^{H}_{ab}(\lambda;x). \tag{2.22}$$

Moreover, it was shown that this additional term is a pure diffeomorphism. Therefore, although the components of the metric in the Ricci flow and the Ricci–DeTurck flow differ, they are in fact geometrically identical, and therefore the flow is independent of the choice of the gauge vector ξ. Furthermore, if the vector ξ is chosen such that $\nabla_{(a}\xi_{b)}$ possesses the same symmetries as the initial metric $g^{(0)}_{ab}$, then the flow (2.22) will preserve these symmetries. From a practical point of view, then, we may regard (2.22) as defining a relaxation scheme for solving 2.9 in vacuum, starting from some initial guess $g^{(0)}_{ab}$.[1] The most straightforward way to proceed is via the *method of lines*, where we apply some spatial discretisation scheme to (2.22) to obtain a system consisting of a large number of coupled, nonlinear ODEs. We can then use any of the numerous time stepping methods to evolve the flow until a steady state is reached. A clear advantage of this method is that, with the choice of an explicit time stepper, it is extremely simple to implement and has a very small memory requirement. However, the convergence of an explicit time integrator is conditional on the step size $\delta\lambda$ being below the *Courant–Friedrichs–Lewy (CFL) bound*. For a parabolic system, this bound takes the form $\delta\lambda = \mathcal{O}(\delta x^2)$, making it very expensive to increase spatial resolution. One could contemplate using an implicit integrator instead in order to sidestep this problem, however this requires solving yet another nonlinear system numerically in order to take a time step. At which point, we may as well just use Newton's method to solve the original equations in the first place.

Another problem inherent to the relaxation method is that it is highly susceptible to the presence of a negative mode in the elliptic operator. In our case, if the linearisation of the Ricci tensor, known as the *Lichnerowicz operator*, has negative modes at the stationary solution of interest, then eventually the Ricci flow would start to diverge away from it. For geometries with a single negative mode, it may be possible to prolong the time until divergence by varying the initial guess, but

[1] We can generalise this method to a non-vacuum spacetime by replacing the RHS of (2.22) by the difference between the LHS and RHS of 2.9.

in general this limitation makes Ricci flow a less popular method in the community. It is worth noting, however, that the Ricci flow method does have the appealing theoretical property of being a geometric flow. The failure of the Ricci flow to reach a steady state therefore indicates that the initial guess does not lie in the basin of attraction of a fixed point, irrespective of the gauge choice, whereas the failure of Newton's method to converge could just be the result of a bad gauge choice.

We now move on to discuss Newton's method. Here, let us consider the general problem of finding a solution vector X_i^* to a system of nonlinear equations $F_i(X) = 0$. In our case, X correspond to the metric components, F correspond to the Ricci–DeTurck components, and i, j range over both the tensor indices and the grid points. In other words, we treat each tensor component on each grid point as separate variables. Let us now assume that we have some candidate numerical solution $X^{(0)}$, and that the exact solution X^* occurs at a small ΔX away from it. We can then expand our objective function as a Taylor series

$$0 = F_i(X^*) = F_i(X^{(0)} + \Delta X) = F_i(X^{(0)}) + \left.\frac{\partial F_i}{\partial X_j}\right|_{X^{(0)}} \Delta X_j + \mathcal{O}\Big(\|\Delta X\|^2 \Big). \tag{2.23}$$

Therefore, we can obtain a better candidate solution $X^{(1)}$ via

$$X_i^* \approx X_i^{(1)} := X_i - \left(\left.\frac{\partial F}{\partial X}\right|_{X^{(0)}} \right)^{-1}_{ij} F_j(X^{(0)}). \tag{2.24}$$

For a linear objective function F, this procedure yields an exact solution. For a nonlinear function, this defines a scheme in which the numerical solution can be iteratively refined until the residual is smaller than some set tolerance. The main advantage of Newton's method is that the convergence is quadratic in the number of iterations. However, our derivation above assumes that the candidate function is already within some small neighbourhood of the true solution, and most of the time this is clearly not the case in practice. When the guess $X^{(0)}$ is far from the true solution, the scheme (2.24) would quickly diverge after only a few iterations. In order to control this divergence, we must suppress the size of the correction via

$$X_i^{(1)} := X_i - \lambda \left(\left.\frac{\partial F}{\partial X}\right|_{X^{(0)}} \right)^{-1}_{ij} F_j(X^{(0)}) \quad , \quad 0 < \lambda \leq 1. \tag{2.25}$$

This is referred to as the *Newton line search* method. The most basic form of Newton line search is where λ is fixed to some small value initially, and is gradually increased after some fixed number

of steps. The drawback, however, is that if λ is increased too quickly then the solver would again diverge, whereas if λ is kept small for longer than necessary then we would not benefit from the fast quadratic convergence. An alternative method to control step size is through the use of *backtracking*. In this scheme, the residual of $X^{(1)}$ is first evaluated at a few different values of λ. We can then fit a polynomial on to these sampled residuals to estimate the largest value of λ at which the residual still decreases and apply a step of that size. This usually results in a highly efficient solver, however it is important to note that backtracking assumes a monotonic decrease in residual towards the solution. When this is not the case, the scheme would fail to make progress, and one must then revert to the fixed-step scheme until the solution moves onto more favourable terrain.

The disadvantages of Newton's method is the increased usage of resources in forming the Jacobian matrix $\partial F_i \big/ \partial X_j$ and in solving a large linear system. Furthermore, the construction of the Jacobian matrix is a much more involved programming task, compared to the Ricci flow method. The most straightforward way is to approximate the Jacobian numerically by perturbing each component in the solution vector and calculate the difference in F. However, the naive implementation of this method is prohibitively expensive as it involves calculating F twice for every single element in the Jacobian. When using finite differences, however, we can deduce the sparsity pattern of the Jacobian matrix from the extent of the stencils used, thus significantly reducing the time required to construct the Jacobian. Alternatively, we can evaluate the analytic expression of the Jacobian components for the given family of metrics. Generically, this results in highly complicated expressions, and a computer algebra system such as MATHEMATICA is typically used to generate these. The complexity of these expressions may sometimes also result in a loss of numerical precision.

2.3.3 Convergence testing

A numerical solution x^* to any equation $f(x) = 0$ is never exact, and the residual $f(x^*)$ will always be nonzero. Although a large residual certainly is a bad sign, a small residual is in itself not sufficient evidence that we have obtained a correct solution. In order to ascertain whether our solution can be regarded as a good approximation to the actual analytic solution, we must perform a *convergence test*, whereby we increase the spatial resolution used to construct the solution and study

how the error behaves as the resolution changes. Ideally, we should require that the error decays uniformly throughout the domain, at a rate which is appropriate to the method used (e.g. correct power in the case of finite differences, or exponential in the case of pseudospectral). In practice, however, limited numerical precision on a computer means that convergence will eventually stop when errors become sufficiently small. Nevertheless, we should still expect the correct convergence rate prior to this point. Likewise, strong gradients near boundaries may cause local errors in these regions to be larger, and limited precision may cause convergence in these regions to stop at a lower resolution. It may therefore be sensible to monitor the Euclidean norm of the errors or the median error value, rather than the L_∞ norm which singles out the maximum error from the entire domain.

In the DeTurck method, the solution is an Einstein metric only if the vector ξ vanishes everywhere. Therefore, this vector provides us with a natural measure of error in our numerical solution, and is therefore a good target on which to perform a convergence test. In each of our projects, we construct a number of numerical solutions to cover a range of parameters. It can be extremely expensive to perform a convergence test on every individual solution in these circumstances. In order to keep resource usage to a reasonable level, we choose two or three solutions on which to perform convergence tests. If the rest of the solutions interpolate smoothly between these representative, tested solutions, then we can be fairly certain that our entire solution set can be trusted.

CHAPTER 3

CFTs on a Rotating Black Hole Background

The contents of this chapter was published as [55] in collaboration with Pau Figueras. All computer code, numerical computation, and analysis, other than section 3.3.4, is entirely my own work. As this work neared completion, we became aware of another independent effort [56], whose initial preprint appeared simultaneously with ours on arXiv. The results of that work agree with those presented in this chapter.

3.1 Introduction

One of the most remarkable effects in quantum gravity is Hawking's realisation that black holes evaporate [57]. Hawking's result motivated the study of quantum fields in black hole backgrounds (discussed in greater detail in [58]), but most of these calculations assume that the fields are non-interacting. Adding interactions is very difficult using this approach. However, in AdS/CFT we can study strongly interacting quantum fields in curved black hole backgrounds, by constructing classical solutions to the bulk Einstein equations with a negative cosmological constant.

The authors of [59, 60, 61] pioneered the study of Hawking radiation in large N_c strongly-coupled field theories which admit gravity duals. As these references point out, the vacuum state in which the CFT is defined is determined by the boundary conditions the geometry satisfies deep in the IR. Then, in principle, the AdS/CFT dictionary enables the extraction of all the physical quantities of the CFT in that particular black hole background and in the chosen vacuum state.

Using numerical methods developed in [45], the authors of [48] constructed an Einstein metric in 5D with a negative cosmological constant, such that the boundary geometry is the 4D asymptotically flat Schwarzschild solution. In the IR the geometry reduces to the Poincaré horizon of AdS_5 and therefore, by the AdS/CFT correspondence, this gravitational solution should describe the large N_c, strong coupling limit of $\mathcal{N} = 4$ super Yang-Mills (SYM) on the background of the Schwarzschild black hole in the Unruh state. Although [48] did not distinguish between the Unruh and Boulware vacua, [62] later pointed out that the Boulware state requires a minimum energy configuration and hence an extremal horizon. This solution enables one to extract the leading $O(N_c^2)$ piece of the vacuum expectation value (vev) of the dual CFT's stress tensor in this particular vacuum state. The results of [48] showed that, remarkably, $\langle T_{ij} \rangle$ is static and is regular on both the future and past event horizons. Furthermore, there is no energy flux at infinity. These results contradict the free field theory expectations (see, for example, [63]), and therefore show that strong interactions can significantly alter the behaviour of quantum fields, especially in curved spaces. Ref. [64] subsequently addressed a closely related problem, namely the construction of static black holes localised on the brane in the single infinite Randall–Sundrum II (RS2) braneworld model [65, 66]. The paper successfully constructed braneworld black holes with both small and large radii relative to the AdS radius of the parent space. This contradicted the non-existence conjecture of [67, 68, 69], which again were based on free field theory intuition.

Describing CFTs in black hole backgrounds in the Hartle–Hawking state in terms of their gravity duals requires different boundary conditions in the IR. In the Hartle–Hawking state, the stress tensor of the plasma at infinity should approach that of a thermal fluid in equilibrium with the boundary black hole. From the bulk perspective, deep in IR the geometry should asymptote to the planar black hole. Ref. [59] further predicted the existence of two families of solutions: *black funnels*, where bulk horizon has a single connected component, and *black droplets*, where the bulk horizon has two disconnected components. The black funnels describe a situation in which the plasma couples strongly to the boundary black hole, resulting in an efficient exchange of heat between them. By contrast, the black droplets describe the weakly coupled situation. These two phases should be connected through a Gregory–Laflamme type transition as the neck of the funnel becomes thinner. Black funnels have been subsequently constructed in [62], while black droplets have been recently constructed in [70].

Black funnels have motivated further research into heat flow in AdS/CFT. Reference [71] showed that rotating BTZ black holes in AdS_3 can be reinterpreted, by means of a change of conformal frame, as describing heat flow in the CFT. However, the conformal symmetry in two dimensions implies that the left and right temperatures of the CFT must be constants, and the bulk horizons are Killing. More recently, [53] considered a stationary funnel solution in global AdS_4, consisting of a single bulk horizon connecting two boundary black holes. The latter have different temperatures, which implies there is a net transport of heat along the bulk horizon, which cannot be Killing. Black holes with non-Killing horizons describing stationary plasma quenches were constructed in [52].

In this chapter we generalise the solutions of [48] by adding rotation to the boundary black hole. While in principle the Hartle–Hawking state cannot exist in rotating black hole backgrounds with Killing horizons (due to issues with superradiance [63]) we are not aware of any such difficulties affecting the Unruh state. Extensive literature exists on the study of *free* quantum fields in static black hole backgrounds, for example in [58] and the references therein. Far less is known about the behaviour of quantum fields in rotating black hole backgrounds. In this chapter, we apply AdS/CFT to compute the stress tensor vev of a strongly interacting 5D quantum CFT in the background of a rotating black hole. Our boundary conditions are such that the vacuum state of the CFT is the Unruh vacuum.

We construct the gravitational dual of a CFT in the background of the 5D Myers–Perry black hole with equal angular momenta [13], as the enhanced symmetry present in this geometry greatly simplifies the numerical computation. However, we expect that our construction should capture most of the qualitative physics of the Unruh state in rotating black hole backgrounds for strongly coupled CFTs that admit gravity duals generally.

The rest of this chapter is organised as follows. In §3.2 we explain our setup and provide details of the numerical construction of our solutions. In §3.3 we present our results, covering both geometrical aspects of our solutions, and the resulting holographic stress tensor. We also show that the falloff behaviour of our numerical stress tensor is consistent with that of the linearised gravitational field on the brane in the Randall–Sundrum (RS2) model. This is obtained by generalising the calculation of [72] to an arbitrary number of dimensions.

3.2 Rotating boundary black holes in Poincaré AdS

In this section, we describe our construction of rotating black holes in Poincaré AdS. Specifically, we seek Einstein metrics in 6D with a negative cosmological constant, whose boundary metrics belong to the same conformal class as the 5D Myers–Perry (MP) black hole [13]. Away from the horizon, the spacetime geometry approaches the Poincaré horizon of AdS_6. Through AdS/CFT, our solution is dual to a CFT living on the background of the MP black hole. As discussed in [48, 62], our boundary conditions in the IR should imply that the dual CFT is in the Unruh vacuum state.

Since the boundary horizon is rotating and Killing, the bulk horizon will have the same angular velocity as the boundary black hole by continuity. Therefore, our solutions are the first examples of rotating horizons in the Poincaré patch of AdS. (The planar limit of a rotating black hole in global AdS yields a boosted black brane.) We restrict ourselves to the case of *non-extremal* black holes, and the temperatures of the boundary and bulk horizons are the same.

In four spatial dimensions, the black hole can have two independent rotations. In the general case, the stationary MP spacetime possesses a $\mathbb{R}_t \times U(1) \times U(1)$ symmetry. However, when the two angular velocities (and hence angular momenta) are equal, the symmetry is enhanced to $\mathbb{R}_t \times SU(2)$, and the metric depends only on the radial coordinate. Therefore, if we take our boundary MP black hole to have equal angular momenta, then the full bulk metric depends only on two coordinates. This greatly reduces the cost of our numerical computation, as we need only solve the PDE in two dimensions. The 5D MP solution with equal angular momenta shares a number of properties with the 4D Kerr solution, such as a smooth extremal limit, and the absence of an ultraspinning regime [73]. We therefore expect that the physics of $\mathcal{N} = 4$ SYM on the Kerr black hole background should be qualitatively similar to that described in this chapter. Indeed, we believe our solution should capture the essential physics of quantum field theories on rotating black hole backgrounds generally.

3.2.1 Equal angular momenta 5D Myers–Perry black hole

We begin this subsection by reviewing the geometry of the 5D Myers–Perry black hole with equal angular momenta. We then introduce new coordinates which will be useful for our subsequent numerical construction of the gravitational duals of CFTs in rotating black hole backgrounds.

The 5D Myers–Perry black hole with equal angular momenta is described by the following line element [13]:

$$\mathrm{d}s^2_{\mathrm{MP}_5} = -\mathrm{d}t^2 + \left[\frac{R^2(R^2+a^2)}{(R^2+a^2)^2-\mu R^2}\right]\mathrm{d}R^2 + \left[\frac{\mu}{R^2+a^2}\right]\left(\mathrm{d}t + \frac{a}{2}\sigma^3\right)^2 + \left(R^2+a^2\right)\mathrm{d}\Omega^2_{(3)}\,, \tag{3.1}$$

where μ and a are the mass and angular momentum parameters respectively, and

$$\begin{aligned}\sigma^1 &= -\sin\psi\,\mathrm{d}\theta + \sin\theta\cos\psi\,\mathrm{d}\phi\\ \sigma^2 &= \cos\psi\,\mathrm{d}\theta + \sin\theta\sin\psi\,\mathrm{d}\phi\\ \sigma^3 &= \cos\theta\,\mathrm{d}\phi + \mathrm{d}\psi\end{aligned} \tag{3.2}$$

are the left-invariant one-forms on $SU(2)$. Here $\mathrm{d}\Omega^2_{(3)} = \frac{1}{4}\left((\sigma^1)^2+(\sigma^2)^2+(\sigma^3)^2\right)$ is the round metric on the unit 3-sphere. The event horizon is located at $R = R_H$, where R_H is the largest real root of the equation $(R_H^2+a^2)^2 - \mu R_H^2 = 0$. In the following, it will be useful to express the mass parameter μ in terms of R_H and a as

$$\mu = \frac{(R_H^2+a^2)^2}{R_H^2}\,. \tag{3.3}$$

It is then straightforward to find a convenient compact radial coordinate by defining

$$R^2 + a^2 = \frac{R_H^2+a^2}{(1-r^2)^2}\,, \tag{3.4}$$

so that $r = 0$ corresponds to the event horizon (or, more precisely, the bifurcation surface) and $r \to 1$ corresponds to spacelike infinity. In terms of the new radial coordinate r, we can rewrite

(3.1) as

$$\mathrm{d}s^2_{\mathrm{MP}_5} = -r^2 T_0(r)\,\mathrm{d}t^2 + \frac{4R_0(r)}{\left(1-r^2\right)^4}\,\mathrm{d}r^2 + \frac{B_0(r)}{4\left(1-r^2\right)^2}\left(\sigma^3 - \Omega_0(r)\mathrm{d}t\right)^2 + \frac{S_0}{4\left(1-r^2\right)^2}\,\mathrm{d}\Omega^2_{(2)}\,, \tag{3.5}$$

where

$$\begin{aligned} R_0(r) &= \frac{R_H^2\left(R_H^2+a^2\right)}{\left(2-r^2\right)\left(R_H^2 - a^2(1-r^2)^2\right)} \\ S_0(r) &= R_H^2 + a^2 \\ T_0(r) &= \frac{\left(2-r^2\right)\left(R_H^2 - a^2(1-r^2)^2\right)}{R_H^2 + a^2\left(1-r^2\right)^4} \\ B_0(r) &= \frac{\left(R_H^2+a^2\right)\left(R_H^2 + a^2(1-r^2)^4\right)}{R_H^2} \\ \Omega_0(r) &= -\frac{2a\left(1-r^2\right)^4}{R_H^2 + a^2\left(1-r^2\right)^4}, \end{aligned} \tag{3.6}$$

and $\mathrm{d}\Omega^2_{(2)} = (\sigma^1)^2 + (\sigma^2)^2$ is the metric on the round unit 2-sphere. Note that, when written in this form, the metric exhibits the structure of a $U(1)$ fibration over an S^2 base. As we will see in §3.2.2, we can consistently construct an extension of this metric into the bulk of AdS_6 while preserving this structure. In these coordinates, the null generator of the future event horizon is given by,

$$k = \frac{\partial}{\partial t} + \Omega_H \frac{\partial}{\partial \psi}, \tag{3.7}$$

and the angular velocity and surface gravity of the horizon are

$$\Omega_H = \Omega_0(0) = \frac{2a}{R_H^2 + a^2}\,, \qquad \kappa^2 = \frac{T_0(0)}{4R_0(0)} = \frac{\left(R_H^2 - a^2\right)^2}{R_H^2\left(R_H^2+a^2\right)^2}\,. \tag{3.8}$$

The form of the metric in (3.5) is particularly useful because the boundary conditions at the horizon are straightforward, as will be discussed in the following subsection. In fact, since any rotating black hole solution with a Killing horizon and with an orthogonally transitive isometry group can be cast in an analogous form, the boundary conditions we discuss below are general. Alternatively, we could have chosen to work in co-rotating coordinates as in [46] at the expense of losing the manifest asymptotic flatness of the metric. This could potentially be problematic when extending the metric into the AdS bulk if we wish to ensure that the spacetime asymptotes to the Poincaré horizon of AdS_6 far from the horizon.

3.2.2 Metric ansatz

Let us first consider the metric on the Poincaré patch of AdS_6:

$$ds^2_{AdS_6} = \frac{\ell^2}{z^2}(dz^2 - dt^2 + dR^2 + R^2 d\Omega^2_{(3)}) \,. \tag{3.9}$$

For our construction it is convenient to introduce new coordinates that are adapted to the Poincaré horizon of AdS. To do so, we generalise the coordinate change introduced in [48] and define new coordinates (r, x) as

$$z = \frac{(1-x^2)\sqrt{R_H^2 + a^2}}{1-r^2} \,, \qquad R = \frac{x\sqrt{(2-x^2)(R_H^2 + a^2)}}{1-r^2} \,, \tag{3.10}$$

in terms of which (3.9) becomes

$$ds^2_{AdS_6} = \frac{\ell^2(1-r^2)^2}{(1-x^2)^2}\left[\frac{4\,dx^2}{(2-x^2)(1-r^2)^2} - \frac{dt^2}{R_H^2 + a^2} + \frac{4r^2\,dr^2}{(1-r^2)^4} + \frac{x^2(2-x^2)}{(1-r^2)^2}\,d\Omega^2_{(3)}\right] . \tag{3.11}$$

In these coordinates the conformal boundary of AdS is in the asymptotic region $x \to 1$, and $x = 0$ is the fixed point set of the $SO(4)$ symmetry. These coordinates are adapted to the Poincaré horizon of AdS in the sense that the latter lies at $r = 1$ and $x < 1$. The point $x \to 1$ and $r \to 1$ corresponds to spacelike infinity at the boundary.

We wish to construct a 6-dimensional, asymptotically locally AdS spacetime, with a boundary metric in the same conformal class as the 5D Myers–Perry black hole, and which approaches the Poincaré horizon of AdS far from the conformal boundary. The isometry group of the boundary metric is $\mathbb{R}_t \times SU(2) \times U(1)$. In our construction, we impose this isometry on the entire bulk geometry. It has been shown in 4D that regular and *static* asymptotically locally AdS metrics inherit the isometry group of the boundary metric [74]. If the same result holds in 6D and in the presence of rotation, then our solutions would be the only ones which are compatible with the Myers–Perry boundary condition. As such, our full spacetime metric should have the full $\mathbb{R}_t \times SU(2) \times U(1)$ isometry group of the boundary metric. We restrict ourselves to the class of metrics which are closed under diffeomorphisms that preserve these symmetries. The general line

element belonging to this class is given by

$$\begin{aligned} \mathrm{d}s^2 = \frac{\left(1-r^2\right)^2}{\left(1-x^2\right)^2}\Bigg[& \frac{4\,X(x,r)}{\left(2-x^2\right)\left(1-r^2\right)^2}\,\mathrm{d}x^2 + \frac{2\,x\,r\left(R_H^2+a^2\right)Y(x,r)}{\left(1-r^2\right)^3}\,\mathrm{d}x\,\mathrm{d}r \\ & - r^2\,T_0(r)\,T(x,r)\,\mathrm{d}t^2 + \frac{x^2\left(2-x^2\right)B_0(r)\,B(x,r)}{4\left(1-r^2\right)^2}\left(\sigma_3-\Omega(x,r)\mathrm{d}t\right)^2 \\ & + \frac{x^2\left(2-x^2\right)S_0(r)\,S(x,r)}{4\left(1-r^2\right)^2}\,\mathrm{d}\Omega_{(2)}^2 + \frac{4\,R_0(r)\,R(x,r)}{\left(1-r^2\right)^4}\,\mathrm{d}r^2\Bigg]\,, \end{aligned} \tag{3.12}$$

where $\{X(x,r), Y(x,r), T(x,r), B(x,r), \Omega(x,r), R(x,r), S(x,r)\}$ are the unknown functions to be solved for, and $\{T_0(r), R_0(r), B_0(r), S_0(r)\}$ are the Myers–Perry functions given in (3.6). As above, the coordinate ranges are $0 \le x < 1$ and $0 \le r < 1$, but we will effectively treat the asymptotic regions (the conformal boundary of AdS $x = 1$, and the Poincaré horizon $r = 1$) as boundaries. We will therefore work on the square domain $0 \le x \le 1$ and $0 \le r \le 1$. On this domain we assume that all our unknown functions are smooth and bounded. Note that because we have chosen coordinates which make the isometries of the metric manifest, $r = 0$ (the horizon) and $x = 0$ (the fixed point set of the $SO(3)$ symmetry) appear as boundaries of our domain. However, as discussed in [48, 46], some of these boundaries are in fact fictitious after suitable smoothness conditions on the metric functions are imposed.

3.2.3 Boundary conditions

Comparing (3.12) with (3.5), we see that at $x = 1$ we must impose the following Dirichlet boundary conditions:

$$\begin{aligned} & X(1,r) = \ell^2\,, \quad Y(1,r) = 0\,, \quad \Omega(1,r) = \Omega_0(r)\,, \\ & T(1,r) = R(1,r) = B(1,r) = S(1,r) = \frac{\ell^2}{R_H^2+a^2}\,. \end{aligned} \tag{3.13}$$

This choice ensures that as $x \to 1$ the metric becomes

$$\mathrm{d}s^2 \sim \frac{\ell^2\left(1-r^2\right)^2}{\left(1-x^2\right)^2}\left[\frac{4}{\left(2-x^2\right)\left(1-r^2\right)^2}\,\mathrm{d}x^2 + \frac{1}{\left(R_H^2+a^2\right)}\,\mathrm{d}s^2_{\mathrm{MP}_5}\right]\,, \tag{3.14}$$

which shows that the boundary metric lies in the same conformal class as the 5D Myers–Perry black hole. It is worth noting that we allow for a constant factor of $\left(R_H^2+a^2\right)^{-1}$ multiplying

$ds^2_{\text{MP}_5}$, so that the boundary conditions at $x = 1$ and $r = 1$ are compatible for any choice of R_H and a. This is necessary to ensure smoothness of the unknown functions.

At $r \to 1$, we want our metric to smoothly approach the Poincaré horizon of AdS_6. According to the discussion of boundary conditions for extremal horizons presented in [48], we must therefore impose

$$
\begin{aligned}
&X = \ell^2 + (1-r)\, X_1\,, \quad Y = (1-r)\, Y_1\,, \quad T = \frac{\ell^2}{R_H^2 + a^2} + (1-r)\, T_1\,,\\
&R = \frac{\ell^2}{R_H^2 + a^2} + (1-r)\, R_1\,, \quad B = \frac{\ell^2}{R_H^2 + a^2} + (1-r)\, B_1\,, \quad S = \frac{\ell^2}{R_H^2 + a^2} + (1-r)\, S_1\,,\\
&\Omega = (1-r)\Omega_1\,, \quad T_1|_{r=1} - R_1|_{r=1} = \text{constant}
\end{aligned}
\tag{3.15}
$$

where all functions with the subscript $_1$ are smooth at $r = 1$ and at $x < 1$. Note that $R_0(1) = B_0(1) = S_0(1) = R_H^2 + a^2$ and also that $\mathrm{d}\Omega^2_{(3)} = \frac{1}{4}\left(\mathrm{d}\Omega^2_{(2)} + \left(\sigma^3\right)^2\right)$. As shown in [48], these boundary conditions are necessary and sufficient for our spacetime to asymptote to the Poincaré horizon of AdS_6. We do not in fact impose the condition $T_1|_{r=1} - R_1|_{r=1} = \text{constant}$, but we do check that it is satisfied by our solutions as consequence of the equations of motion. Hence, we can use this condition as an estimate of the numerical error.

At $x = 0$ both the S^2 of the base space and the $U(1)$ fiber degenerate. To ensure that they do so smoothly, the functions X, Y, T, R, B, Ω, S must be smooth functions of x^2 near $x = 0$ and hence they should obey a Neumann boundary condition there. In addition, to avoid the presence of conical singularities at $x = 0$ we must require

$$
\frac{B_0(r)\, B(0,r)}{X(0,r)} = \frac{S_0(r)\, S(0,r)}{X(0,r)} = 1 \quad \text{at} \quad x = 0\,. \tag{3.16}
$$

Smoothness of the manifold at the horizon requires all metric functions to be smooth in r^2 near $r = 0$; hence they must obey a Neumann boundary condition there. In addition, as discussed in [46], we must ensure that our metric has the same surface gravity and angular velocity as the reference metric (described below), which can be achieved imposing

$$
T(x,0) = R(x,0)\,, \qquad \Omega(x,0) = \Omega_0(0)\,. \tag{3.17}
$$

Then the surface gravity and angular velocity of our spacetime will be given by (3.8).

In order to solve the Einstein–DeTurck equations (2.8) we have to specify a reference metric $\bar{g}$ on the same manifold $\mathcal{M}$ as the physical spacetime metric g that we are seeking. This is equivalent to choosing a particular gauge. In our case, we have chosen $\bar{g}$ to be given by (3.12) with

$$\begin{aligned} &X = \ell^2\,, \quad Y = 0\,, \quad T = R = S = \frac{\ell^2}{R_H^2 + a^2}\,, \\ &\Omega = \Omega_0\,, B = \frac{\ell^2}{R_H^2 + a^2}\left[x^2 + \frac{R_H^2(1-x^2)}{B_0(R_H^2 + a^2)}\right]. \end{aligned} \tag{3.18}$$

Finally, note that the boundary conditions above are compatible with $\xi^a|_{\partial\mathcal{M}} = 0$. Given our boundary conditions, we cannot rule out *a priori* the existence of 'Ricci solitons' (i.e. solutions to (2.8) with $\xi^a \neq 0$). However, because the problem is elliptic, solutions should be locally unique. Therefore an Einstein metric cannot be arbitrarily close (in the space of solutions) to a Ricci soliton [45, 48, 46]. As such, we can always check *a posteriori* for a given solution whether $\xi^a = 0$. For all the solutions presented in §3.3 we find that $\xi^a \to 0$ in the continuum limit.

3.2.4 Details of the numerics

For this work, we solve (2.8) by simulating the associated parabolic Ricci–DeTurck flow, as discussed in Section 2.3.2. In order to carry out the flows we must provide an initial metric $g_{ab}|_{\lambda=0}$, and for the results presented in §3.3 we set the initial metric equal to the reference metric. The results should not be qualitatively different for other choices of initial data, but we did not investigate this.

As discussed in [75, 45], convergence of the flow toward the fixed point depends on the stability of the fixed point, which is in turn determined by the presence or absence of negative modes in the spectrum of the linearisation of (2.8) computed around the fixed point. This coincides with the Lichnerowicz operator when the fixed point is an Einstein metric. In this discussion, we assume that this is the case. In our simulations, we found that our initial data always converged to the desired Einstein metric with no fine-tuning. This indicates that our solutions are stable fixed points of the Ricci–DeTurck flow and hence do not posses negative modes. Even though the 5D

Myers–Perry black hole has a negative mode [73] (analogous to the celebrated negative mode of 4D Schwarzschild [76] or Kerr [77]) it seems that pinning the metric to the boundary of AdS projects out this mode. Reference [48] observed the same phenomenon for the 5D bulk case, in which the boundary metric is 4D Schwarzschild.

In AdS/CFT, only the conformal class of boundary metrics matters, and therefore we can choose any representative within the class without loss of generality. For the problem at hand, choosing a representative of the conformal class amounts to fixing the overall length scale of the boundary 5D Myers–Perry black hole. In the results below, we fixed the mass parameter μ to some convenient value and varied the angular momentum parameter a to move along the family of solutions. Note that this implies that both the surface gravity and angular velocity vary along this family.

As in [45], there are two main sources of numerical error. The first is present in any numerical method, and is due to the fact that we discretise the equations (2.22) according to some scheme and seek a solution to the discrete problem. The latter approximates the continuum solution, and, for a consistent discretisation, the error should decrease in the continuum limit at a rate dependent on the discretisation scheme and the degree of differentiability of the solution. The second source of error relates to the fact that it takes infinite Ricci flow time for (2.22) to reach its fixed point, but in practice we can only evolve (2.22) for a finite time. We must therefore ensure that we are sufficiently close to the fixed point that the discretisation error dominates. We do so by monitoring $|\ell^2\, R/30 + 1|_{\max}$ along the flow, which vanishes for an Einstein metric in the continuum limit. We stop the simulation only when this quantity is constant to within some tolerance. The value of the constant provides an estimate of the numerical error of our solution, which decreases as we increase the spatial resolution. Of course, the amount of Ricci flow time that it takes to reach the fixed point depends on how close to extremality the solution is and the resolution of the spatial discretisation.

We simulated the Ricci–DeTurck flow using the method of lines. The two spatial dimensions were discretised onto a square grid of Chebyshev points, and spatial derivatives were approximated by spectral differentiation. The flow was then time-evolved using a third-order Runge–Kutta integrator with a fixed step size. Note also that spatial differentiation is equivalent to simple pointwise multiplication in the spectral space, meaning that we can improve the computational speed by using the discrete cosine transform rather than repeated matrix multiplications.

In our calculation of the curvature tensor, we hard-coded the expression for each of the non-zero components in terms of the metric components and their derivatives. Due to the complexity of this problem, a typical component involves a large number of additions and subtractions of floating point quantities. As the metric becomes singular near the horizon, the relative sizes of these summands can become vastly different. At double-precision, we found that this becomes problematic very close to the fixed point: numerical error begins to dominate and the flow oscillates wildly and does not converge any further. The problem worsens as we increase the number of grid points, since the metric components need to be computed closer to the horizon.

This problem is exacerbated by the fact that, in order to extract the boundary stress tensor, we need to read off the fifth derivative from our numerical data. For each extra derivative that we require, the numerical results must be evaluated more precisely. In some cases, we found that in order to extract the fifth derivative to acceptable accuracy, we had to invoke quadruple-precision floating point arithmetic in our numerical computation. Unfortunately, no current equipment implements this at the hardware level, so we had to rely on software libraries to emulate higher-precision floating point arithmetic, which can be over 20 times slower than native computation.

To optimise the time taken to obtain acceptably accurate solutions, we divided the typical run into three phases. We first evolved the Ricci flow equation numerically using standard 64-bit double-precision floating point arithmetic. Once our error indicators (the Ricci scalar and the norm of DeTurck vector) began to oscillate, we took the numerical solution at the final Ricci flow time and used it as initial data for a second run, carried out using 80-bit "long double" precision. Finally, we take our solution from the second run as initial data for a final run at 128-bit quadruple precision. The entire process typically takes between one and two days. In all cases, the data presented in the next section is the result of the final run at quadruple precision.

In the majority of cases, we found that a 36×36 square Chebyshev grid provided a good balance between accuracy and computational requirements. However, for the slowly-spinning cases, we found that finer grids (45×45) are necessary to maintain acceptable accuracy.

3.3 Numerical results

This section presents our numerical results. To construct our solutions we fixed the mass parameter of the boundary black hole, μ, then varied the angular momentum parameter, a, such that the dimensionless surface gravity is $\kappa\,\mu^{\frac{1}{2}} = \frac{n}{16}$ with $n = 1, \ldots, 16$. Therefore, for our solutions the horizon of the black hole is always non-extremal.

3.3.1 Geometry of the horizon

To visualise the horizon geometry of our solutions as we vary the boundary black hole's angular momentum, we consider the following embedding. We take the geometry of the spatial cross-sections of the horizon transverse to the $U(1)$ fiber,

$$\mathrm{d}s^2_H = \frac{\ell^2}{(1-x^2)^2}\left[\frac{4\,X(x,0)}{2-x^2}\,\mathrm{d}x^2 + \frac{x^2(2-x^2)S_0(0)\,S(x,0)}{4}\,\mathrm{d}\Omega^2_{(2)}\right], \tag{3.19}$$

and embed them into 4D hyperbolic space with the same radius ℓ, $\mathrm{d}s^2_{\mathbb{H}_4} = \frac{\ell^2}{z^2}(dz^2 + \mathrm{d}R^2 + R^2\,\mathrm{d}\Omega^2_{(2)})$, as a curve $R = R(z)$. We then colour-code this curve according to the relative size of the $U(1)$ fibre with respect to the base S^2, which is measured by the function $\alpha(z) = B_0(0)B(x,0)/(S_0(0)S(x,0))$, where the coordinate x is a function of the z coordinate of the ambient hyperbolic space. The freedom in the embedding is fixed by requiring that the radius of the boundary black hole is $R(0) = 1$ in the static case, which fixes the ADM mass of the boundary black hole for any $a \geq 0$. Figure 3.1 presents the results. As this figure demonstrates, the radius of the S^2 becomes smaller as the angular momentum increases. This reflects the fact that the horizon area at fixed mass decreases with greater angular momentum. Note also that the extent of the horizon into the bulk decreases as the angular momentum increases. Moreover, for a fixed value of the angular momentum parameter $a > 0$, the squashing of the horizon three-sphere is more severe at the boundary than in the interior of the spacetime. In fact, the regularity condition (3.16) at the axis of symmetry $x = 0$ implies that the horizon S^3 must be locally round there, ruling out any squashing.

Figure 3.2 depicts the evolution of the curvature invariant $C_{abcd}C^{abcd}\,\ell^4$ along the flow for a

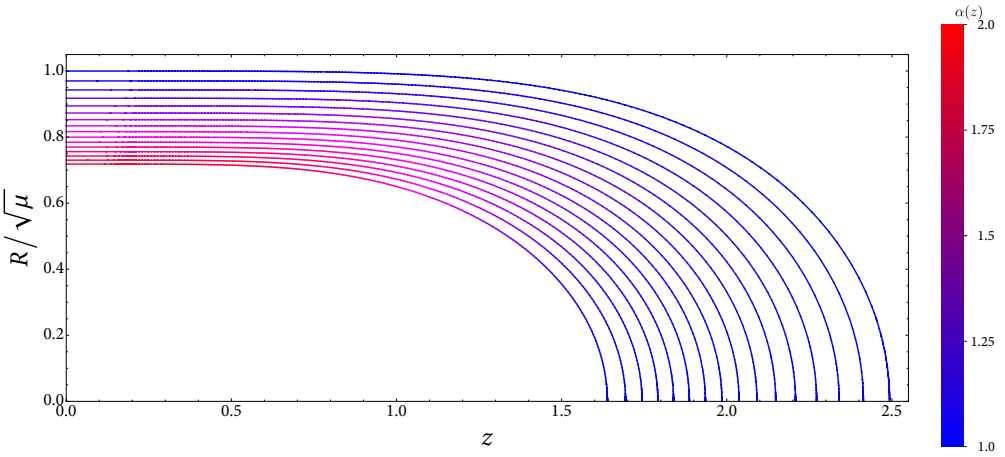

Figure 3.1: Embeddings of the radius $R(z)$ of the horizon S^2 into four-dimensional hyperbolic space, for fixed mass parameter μ of the boundary black hole. At the boundary of hyperbolic space, $z = 0$, the value of this radius function coincides with the corresponding one for the 5D Myers–Perry black hole geometry. The colour code indicates the squashing of the horizon S^3 according to the relative size of the fibre with respect to the base S^2, which is measured by the function $\alpha(z) = B_0(0)B(z,0)/(S_0(0)S(z,0))$. The blue curve with $R(0) = 1$ corresponds to the zero rotation case, and a increases moving inwards in this plot. The squashing becomes more severe near the boundary and with increasing angular momentum of the boundary black hole. Absence of conical singularities at the axis of symmetry $R(z_{\max}) = 0$ implies that there is no squashing there.

typical solution. This is a useful geometric quantity because it provides information about both the curvature of the spacetime and the correctness of our boundary conditions. As the figure shows, $C_{abcd}C^{abcd}\,\ell^4$ never blows up along the flow (and in particular at the fixed point), which suggests that there are no curvature singularities in our domain. In addition, we see that for any λ this curvature invariant vanishes at both the Poincaré horizon of AdS_6 and at the conformal boundary. This confirms that our spacetime has the correct asymptotics near these two boundaries.

3.3.2 Linearised gravity in the higher dimensional RS2 model

From our numerical solutions, we can now proceed to extract the stress tensor of the dual CFT. Firstly, however, it is useful to have a theoretical result to which we can compare our numerical stress tensor. Following the reasoning of [39, 48], the fall-off of the stress tensor in the AdS/CFT solution is given by the correction to the standard gravitational potential due to the presence of a brane. In this section, we therefore generalise the derivation of the linearised gravitational field

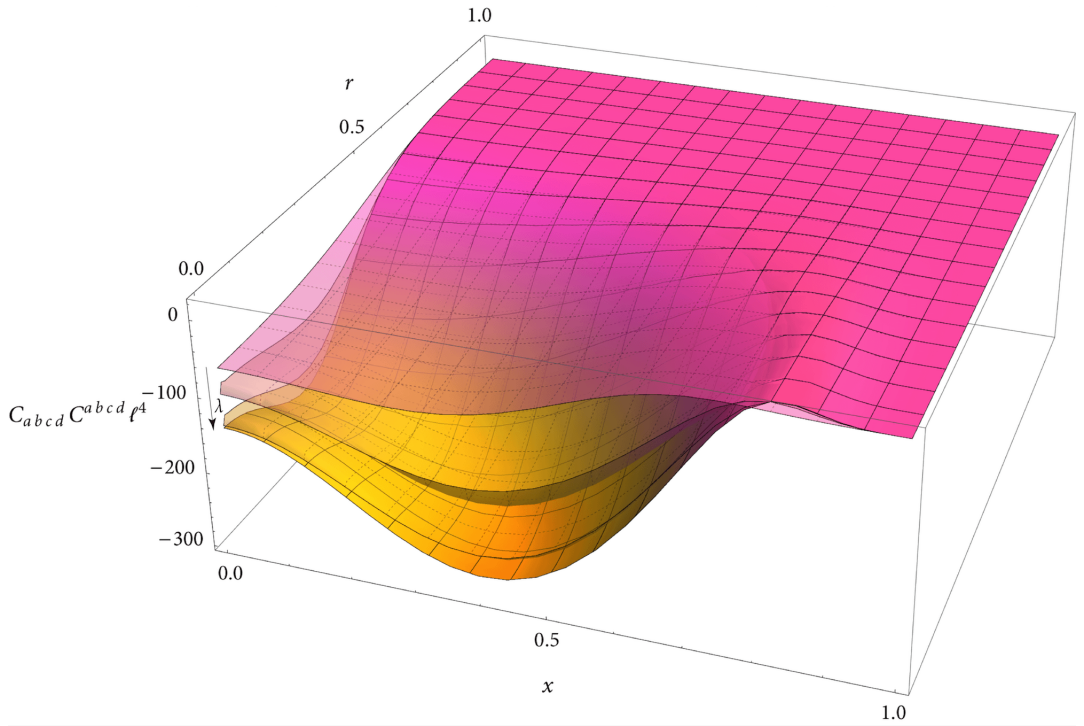

Figure 3.2: $C_{abcd}C^{abcd}\,\ell^4$ along the Ricci flow in the whole domain for the $\kappa\,\mu^{\frac{1}{2}} = 0.15$ solution. The snapshots correspond to the Ricci flow times $\lambda = 0,\ 0.5,\ 1$, and the fixed point (top to bottom in this plot). After a Ricci flow time of $\lambda = 1$ the surface becomes almost indistinguishable from that of the fixed point. As this plot shows, $C_{abcd}C^{abcd}\,\ell^4$ does not blow up anywhere in the domain along the flow. It also vanishes at both the Poincaré horizon and the boundary of AdS, indicating that our spacetime is asymptotically AdS near these boundaries.

in the RS2 braneworld model in [72, 78] to an arbitrary number of spacetime dimensions. For D boundary dimensions, this shows that the energy density of the dual field theory behaves like $1/R^{2D-3}$ in the asymptotic region. In $D = 5$, this gives $1/R^7$ behaviour, which is in agreement with our data, as we show in the next section.

We begin by considering the metric of a D-dimensional brane embedded in a $(D+1)$-dimensional AdS space

$$\mathrm{d}s^2 = \mathrm{d}y^2 + a(y)^2\,\eta_{\mu\nu}\mathrm{d}x^\mu\mathrm{d}x^\nu, \tag{3.20}$$

where $a(y) = \mathrm{e}^{-|y|/\ell}$. For a metric perturbation of the form $\eta_{\mu\nu} \mapsto \eta_{\mu\nu} + h_{\mu\nu}$, the perturbation $h_{\mu\nu}$ satisfies

$$-\frac{1}{2}\left[a(y)^{-2}\,\Box_{(D)}\,h_{\mu\nu} + a(y)^{-(D-2)}\partial_y\Big(a(y)^D\,\partial_y\big(a(y)^{-2}h_{\mu\nu}\big)\Big)\right] = 0\,, \tag{3.21}$$

where we have taken $h_{\mu\nu}$ to satisfy the Randall–Sundrum gauge conditions

$$h_{55} = 0 = h_{\mu 5} \quad , \quad h_{\mu}{}^{\nu}{}_{;\nu} = 0 \quad , \quad h^{\mu}{}_{\mu} = 0. \tag{3.22}$$

We look for normal modes $h_{\mu\nu} = \psi(y)\mathrm{e}^{ik^{\rho}(x-x')_{\rho}}$ with $k^{\rho}k_{\rho} = -m^2$. Imposing the Israel junction condition on the brane (see [72] for details) and then expanding out the operator, we arrive at an equation for $\psi(y)$:

$$\left[a(y)^{-2}m^2 + \partial_y\partial_y - (D-4)\,\ell^{-1}\partial_y - 2\,(D-2)\,\ell^{-2}\right]\psi(y) = -4\ell^{-1}\delta(y)\psi(y). \tag{3.23}$$

Note that in $D = 4$, the coefficient of the first derivative term vanishes, and this expression reduces to the one presented in [72]. The general solution to this ODE is given in terms of exponentially-modulated Bessel functions. As the equation is second-order, there are two coefficients to be fixed for each m: an overall multiplicative constant, which we denote C_m, and a relative coefficient, which we denote k_m. We will now show that latter can be fixed by the aforementioned junction condition, while the former can be fixed by an off-brane boundary condition.

When D is odd, the most convenient form of the general solution is

$$\psi_m(y) = C_m\, a(y)^{-(D-4)/2}\left[J_{-D/2}(m\ell/a(y)) + k_m J_{D/2}(m\ell/a(y))\right]. \tag{3.24}$$

The value of k_m is fixed by the presence of the δ-function. To see this, we integrate (3.23) over a small neighbourhood $[-\varepsilon, \varepsilon]$ around zero, then take $\varepsilon \to 0$. Since we want $\psi(y)$ to be even under $y \mapsto -y$, we must require $\psi(y)$ to be continuous at $y = 0$. Therefore, the δ-function imposes the following *jump condition* on the first derivative of $\psi(y)$:

$$\psi'\left(0^+\right) - \psi'(0^-) = -4\ell^{-1}\psi(0). \tag{3.25}$$

Equation (3.24) then implies that

$$2m\left(J'_{-D/2}(m\ell) + k_m J'_{D/2}(m\ell)\right) = D\,\ell^{-1}\left(J_{-D/2}(m\ell) + k_m J_{D/2}(m\ell)\right), \tag{3.26}$$

and thus

$$k_m = -\left(\frac{2mJ'_{-D/2}(m\ell) + D\ell^{-1}J_{-D/2}(m\ell)}{2mJ'_{D/2}(m\ell) + D\ell^{-1}J_{D/2}(m\ell)}\right). \tag{3.27}$$

In the following we will also need the expansion of $J_\nu(x)$ about zero. For completeness we record it below:

$$J_\nu(x) \sim \frac{x^\nu}{2^\nu\Gamma(1+\nu)} - \frac{x^{\nu+2}}{2^{\nu+2}\Gamma(2+\nu)} + \dots \qquad \text{as } x \to 0. \tag{3.28}$$

Inserting this into (3.27) gives

$$k_m \sim \frac{\Gamma(D/2)}{\Gamma(2-D/2)}\left(\frac{2}{m\ell}\right)^{D-2} \qquad \text{as } m\ell \to 0. \tag{3.29}$$

When D is even, the most convenient form of the general solution is

$$\psi_m(y) = C_m\, a(y)^{-(D-4)/2}\left[Y_{D/2}(m\ell/a(y)) + k_m J_{D/2}(m\ell/a(y))\right]. \tag{3.30}$$

The jump condition now reads

$$2m\left(Y'_{D/2}(m\ell) + k_m J'_{D/2}(m\ell)\right) = D\ell^{-1}\left(Y_{D/2}(m\ell) + k_m J_{D/2}(m\ell)\right), \tag{3.31}$$

and thus

$$k_m = -\left(\frac{2mY'_{D/2}(m\ell) + D\ell^{-1}Y_{D/2}(m\ell)}{2mJ'_{D/2}(m\ell) + D\ell^{-1}J_{D/2}(m\ell)}\right). \tag{3.32}$$

We now need the expansion of Y_n about zero for $n \in \mathbb{N}$:

$$Y_n(x) \sim -\frac{2^n\,(n-1)!}{\pi x^n} - \frac{2^{n-2}\,(n-2)!}{\pi x^{n-2}} - \dots \qquad \text{as } x \to 0. \tag{3.33}$$

Therefore, in this case k_m is given asymptotically by

$$k_m \sim \frac{(D/2-1)!\,(D/2-2)!}{\pi}\left(\frac{2}{m\ell}\right)^{D-2} \qquad \text{as } m\ell \to 0. \tag{3.34}$$

We can summarise these results to apply in any number of dimensions D by writing $k_m \sim \frac{\kappa_{(D)}}{(m\ell)^{D-2}}$, where $\kappa_{(D)}$ is a constant that depends only on D.

The overall constant C_m is determined by requiring that the eigenmodes are finite as $m \to 0$ for

any value of $|y|$. We can therefore consider the behaviour of the eigenmodes as we move far away from the brane, taking $|y|$ to be sufficiently large. The large argument asymptotics of the Bessel functions are given by

$$J_\nu(x) \sim \sqrt{\frac{2}{\pi x}} \cos\left(x - \frac{(2\nu+1)\pi}{4}\right) \tag{3.35}$$

$$Y_\nu(x) \sim \sqrt{\frac{2}{\pi x}} \sin\left(x - \frac{(2\nu+1)\pi}{4}\right) \qquad \text{as } x \to \infty. \tag{3.36}$$

For convenience, let us write $\hat{y} \equiv a(y)^{-1} = \mathrm{e}^{|y|/\ell}$. Then, for large $m\ell\hat{y}$, the eigenmodes behave like

$$\psi_m(\hat{y}) \sim C_m \sqrt{\frac{2\hat{y}^{D-5}}{\pi m\ell}} \left[\sin\left(m\ell\hat{y} - \frac{(D+1)\pi}{4}\right) + \frac{\kappa_{(D)}}{(m\ell)^{D-2}} \cos\left(m\ell\hat{y} - \frac{(D+1)\pi}{4}\right)\right]. \tag{3.37}$$

Far from the brane, we expect the modes to behave like plane waves [66], and therefore the small-m dependence of C_m should be

$$C_m \sim (m\ell)^{D-3/2} \qquad \text{as } m\ell \to 0, \tag{3.38}$$

neglecting any numerical factors. Finally, since C_m is a constant independent of y, the above dependence must also hold on the brane $y = 0$. We can therefore read off the small-m dependence of ψ_m as

$$\psi_m \sim (m\ell)^{(D-3)/2}. \tag{3.39}$$

The gravitational potential is essentially given by the Green's function for (3.23), which consists of a superposition of all the eigenfunctions ψ_m that we have just derived. We are only interested in the gravitational field on the brane itself in the far field region, so henceforth we set $y = 0$ and do not keep track of the overall numerical coefficients. Note that there is a continuum of KK eigenmodes as well as a discrete zero mode, thus

$$G(x,x')\big|_{y=0} = -\int \frac{\mathrm{d}^D k}{(2\pi)^D} \mathrm{e}^{ik_\mu (x-x')^\mu} \left(\frac{\ell^{-1}}{|\mathbf{k}|^2 - \omega^2} + \int_0^\infty \mathrm{d}m \frac{\psi_m(0)^2}{|\mathbf{k}|^2 + m^2 - \omega^2}\right). \tag{3.40}$$

Since we are also only interested in the stationary state, we can integrate out the t' dependence,

leaving just

$$G(\mathbf{x},\mathbf{x}')\big|_{y=0} = -\int \frac{\mathrm{d}^{D-1}\mathbf{k}}{(2\pi)^D} \mathrm{e}^{i\mathbf{k}\cdot(\mathbf{x}-\mathbf{x}')} \left(\frac{\ell^{-1}}{|\mathbf{k}|^2} + \int_0^\infty \mathrm{d}m \frac{\psi_m(0)^2}{|\mathbf{k}|^2+m^2} \right). \tag{3.41}$$

The first term gives rise to the usual $1/r^{D-3}$ potential for standard gravity on the brane, while the second term is responsible for the correction due to the KK modes; this is the relevant term for us. Thus, we define

$$V_{KK} = -\int \frac{\mathrm{d}^{D-1}\mathbf{k}}{(2\pi)^D} \mathrm{e}^{i\mathbf{k}\cdot(\mathbf{x}-\mathbf{x}')} \left(\int_0^\infty \mathrm{d}m \frac{\psi_m(0)^2}{|\mathbf{k}|^2+m^2} \right). \tag{3.42}$$

Changing the order of integration and doing first the integral over the momenta $\mathbf{k}$, gives rise to a Yukawa-type potential

$$V_{KK} = \frac{1}{r^{D-3}} \frac{\Gamma(D/2-1/2)}{2(D-3)\pi^{(D-1)/2}} \int_0^\infty \mathrm{d}m\, F_{(D-3)/2}(mr)\psi_m(0)^2\,, \tag{3.43}$$

where we have defined $F_\nu(\xi) \equiv \frac{1}{2^{\nu-1}\Gamma(\nu)}\xi^\nu K_\nu(\xi)$, and K_ν is the modified Bessel function of the second kind. Indeed, for $D=4$ we have precisely $F_{1/2}(mr) = \mathrm{e}^{-mr}$. For a general D, $F_{(D-3)/2}$ generalises this exponential screening, as its asymptotic behaviour is given by

$$F_{(D-3)/2}(mr) \sim \frac{\sqrt{\pi}\,(mr)^{D/2-2}\,\mathrm{e}^{-mr}}{2^{D/2-2}\Gamma(D/2-3/2)} \left[1+\mathcal{O}\big((mr)^{-1}\big)\right]. \tag{3.44}$$

We extract the leading order correction to the linearised gravitational potential by integrating over m. As in [66, 72], we note that in the large-r limit the integral is dominated by the contribution from small-m modes. Using (3.39) in (3.43) yields

$$V_{KK} \sim \frac{1}{r^{2d-5}} \qquad (r \gg 1)\,, \tag{3.45}$$

so that the full gravitational potential on the brane is

$$V(r) \sim \frac{G_N M}{r^{D-3}} \left(1 + \frac{\alpha\,\ell^{D-2}}{r^{D-2}}\right) \tag{3.46}$$

where α is some non-vanishing (dimensional dependent) numerical factor, and G_N is Newton's

constant on the brane. Then the metric perturbation on the brane will be

$$h_{00} \sim V(r)\,. \tag{3.47}$$

The other components of the metric perturbation can be shown to exhibit similar behaviour in the far field region.

Following [39, 64] and considering the limit in which the brane is close to the boundary of AdS, one can use the correction to the Einstein tensor on the brane induced by (3.47) to read off the components of the stress tensor in the far field regime:

$$\delta G_{\mu\nu} = 16\,\pi\,G_N\,\langle T^{CFT}_{\mu\nu}\rangle\,. \tag{3.48}$$

Note that only the term proportional to α in (3.46) contributes to the left hand side of this equation, which implies

$$\langle T^{CFT}_{\mu\nu}\rangle \sim \frac{1}{r^{2D-3}}\,. \tag{3.49}$$

For $D = 5$, this is precisely the same behaviour that our data exhibits.

3.3.3 Boundary stress tensor

Our bulk 5D solutions can be trivially uplifted on half an S^4 to massive type IIA supergravity in 10D, which arises as the low energy limit of intersections of D4 and D8 branes in type I′ string theory [79]. The CFT that arises as the fixed point of this system under RG flow is USp($2N_c$) 5D SYM, with matter consisting of a hypermultiplet in the antisymmetric representation of the gauge group and N_f hypermultiplets in the fundamental [80]. This is the CFT dual to our supergravity solutions. We extract the vacuum expectation value of the stress tensor of this dual CFT from the bulk gravity solution, using the standard holographic renormalisation prescription [40]. For a 5D bulk spacetime, the expectation value of the stress tensor of the dual CFT is given by

$$\langle T_{ij}\rangle = \frac{5\,\ell^4}{16\,\pi\,G_6}\,g^{(5)}_{ij}\,, \tag{3.50}$$

where $g^{(5)}_{ij}$ denotes the 5th order term in the near boundary expansion of the metric in Fefferman–Graham (FG) coordinates, and G_6 is the 5D bulk Newton's constant.

Since we are interested in spacetimes with the isometry group $\mathbb{R}_t \times SU(2) \times U(1)$, we can consistently truncate the general form of the line element in FG coordinates to this particular class of metrics. Solving the Einstein equations in the usual near boundary expansion, we find that the stress tensor is given by

$$\langle T_{ij} \rangle \, \mathrm{d}x^i \mathrm{d}x^j = \frac{5\,\ell^4}{16\,\pi\,G_6} \left[T_5(R)\,\mathrm{d}t^2 + R_5(R)\,\mathrm{d}R^2 + B_5(R)(\sigma^3 - \Omega_5(R)\,\mathrm{d}t)^2 + S_5(R)\,\mathrm{d}\Omega^2_{(2)} \right], \tag{3.51}$$

where T_5, R_5, B_5, Ω_5, S_5 can be extracted from our numerical solutions. These functions are not all independent: the Einstein equations at 5th order in the near boundary expansion in FG coordinates impose an algebraic constraint, which is nothing but the tracelessness condition of the dual stress tensor. With an odd number of boundary dimensions, the stress tensor is always traceless regardless of whether the boundary metric is Ricci flat, because there is no gravitational conformal anomaly. At 6th order one finds a differential constraint which implies that (3.51) is covariantly conserved. Therefore, we find that the stress tensor is fully specified in terms of three independent functions. However, in our construction we will independently extract the five functions from our numerical data, then use the trace and divergence of $\langle T_{ij} \rangle$ to estimate the associated numerical error. We find that the trace is largest near the horizon, with a magnitude of around 10^{-5} to 10^{-3} (the fast rotating cases give smaller values). This quickly decays to well below 10^{-9} in all cases away from the horizon.

To obtain the components of the stress tensor in terms of our numerical solutions we must change from the working coordinates (x, r) used in our ansatz (3.12) into the FG coordinates (z, R) that we used to derive (3.51). We proceed as in [52] and determine the change of coordinates $z = z(x, r)$, $R = R(x, r)$ in a near boundary expansion, requiring that $\xi^a = 0$ at each order in $\left(1 - x^2\right)$. Note that in our setup ξ^a only has two non-vanishing components, so imposing $\xi^a = 0$ order by order determines the coordinate change completely.

The dual stress tensor $\langle T_{ij} \rangle$ is symmetric in its two indices, but because the boundary metric (3.1) is not positive definite, the linear map $\langle T^i{}_j \rangle$ from vectors to vectors need not be diagonalisable. However, in our case it is, and for all values of the angular momentum parameter a that we have

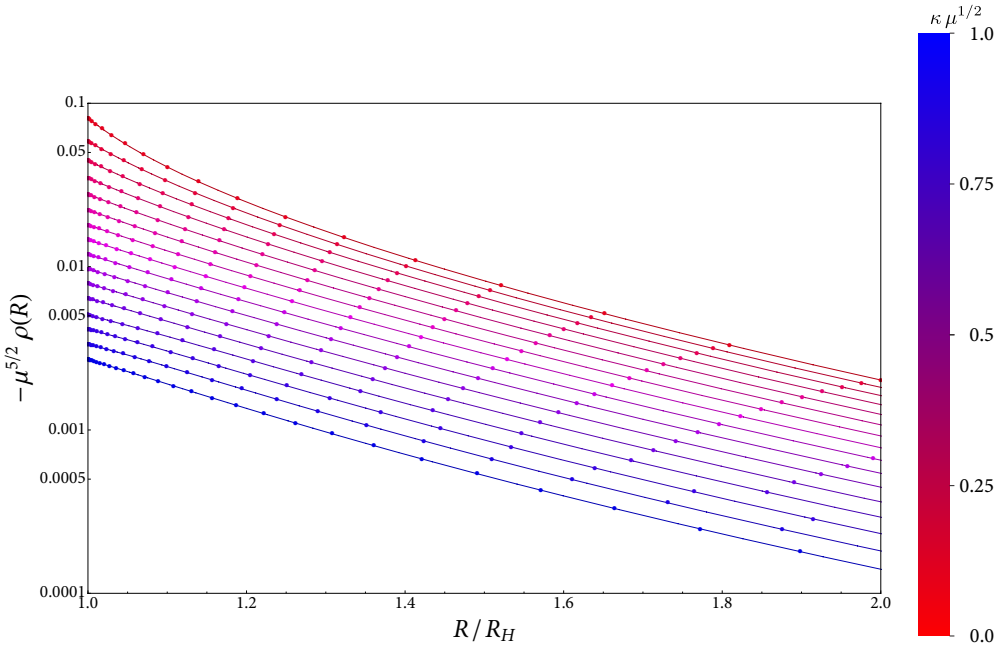

Figure 3.3: A plot of (minus) the energy density, for different values of the rotation parameter versus the Boyer–Lindquist radial coordinate R. We have used a logarithmic scale on the y axis in order to enhance the near horizon region. As this plot shows, the energy density is negative everywhere and it becomes more negative as the rotation of the black hole increases.

considered there is one timelike eigenvector and four spacelike eigenvectors. We can identify the corresponding eigenvalues with the energy density and the pressures in the following way. We write the stress tensor as

$$\langle T^i{}_j \rangle = \rho(R)\, t^i \otimes t_j + p_1(R)\big[(s^1)^i \otimes (s^1)_j + (s^2)^i \otimes (s^2)_j\big] + p_2(R)\, (s^3)^i \otimes (s^3)_j + p_3(R)\, R^i \otimes R_j \tag{3.52}$$

where t^i is the unique timelike eigenvector normalised so that $t^i t_i = -1$, and $(s^\flat)^i$ and R^i are the orthonormal spacelike vectors in the 5D Myers–Perry background[1] written in Boyer–Lindquist-type coordinates (3.1). In this way, we identify $\rho(R)$ as the energy density of the plasma seen by a local observer with velocity t^i, and the $p_i(R)$ as the corresponding pressures.[2]

In Figure 3.3 we plot the energy density for different values of the rotation parameter of the background. First, note that the energy density is *negative* everywhere in our domain for all $a \geq 0$.

[1] We choose $(s^\flat)_i \sim (\sigma^\flat)_i$ with the obvious proportionality factors and $R_i = \sqrt{\frac{R^2(R^2+a^2)}{(R^2+a^2)^2-\mu R^2}}\,(dR)_i$.

[2] Recall that if k^i is a unit timelike and future directed vector field, then $-T^i{}_j\, k^j$ represents the energy density seen by an observer with velocity k^a. So $-\rho(R)$ is the eigenvalue of the matrix $T^i{}_j$.

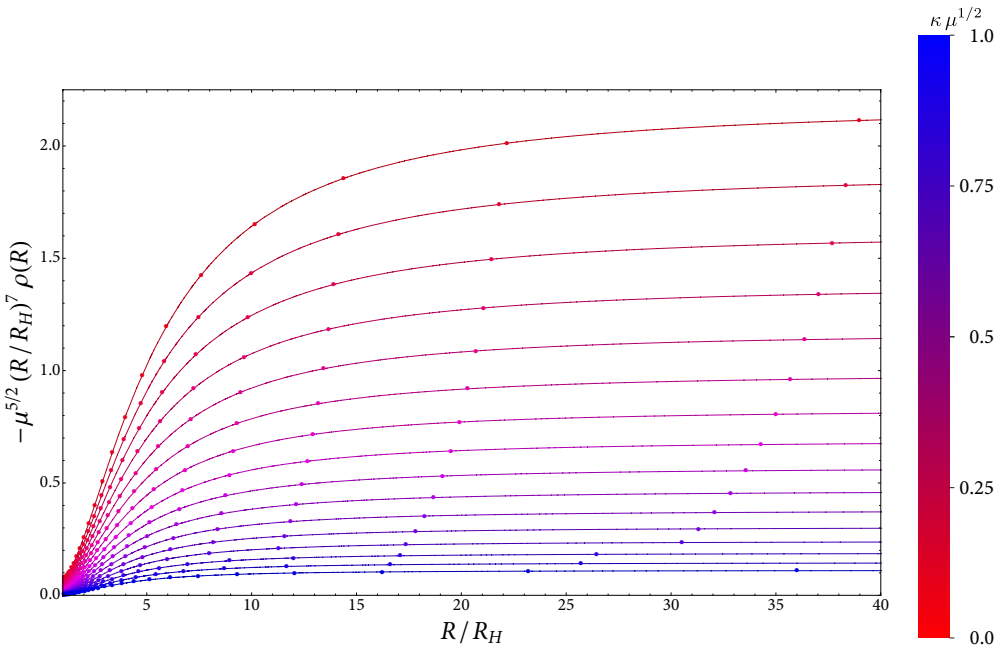

Figure 3.4: A plot of the rescaled energy density for different values of the rotation parameter. This plot shows that the energy density decays like $1/R^7$ in the asymptotic $R \to \infty$ region. The other components of the stress tensor have a similar behaviour near infinity. The colour code is the same as in Figure 3.3.

Moreover, the energy density is a monotonically increasing function of the radial coordinate: it becomes less negative for large R, and has the largest absolute value at the horizon. Furthermore, as the angular momentum increases (moving up the figure) the energy density becomes more negative at the horizon, but it is always finite. By continuity, the energy density should also be finite at extremality.

Next we consider the pressures. As in [52] we find that the radial pressure is positive, while the pressures along the angular directions are negative. There is therefore an anisotropy in the system. Note that this anisotropy is already present in the static limit. Adding rotation to the boundary black hole breaks the full $SO(4)$ spherical symmetry of the static case down to $SU(2) \times U(1)$, which results in some further anisotropy between the fibre and the sphere directions. This simply reflects the squashing of the S^3 of the boundary black hole. As the angular momentum of the boundary black hole increases, the horizon S^3 becomes more squashed, which gives rise to a more negative pressure along the fibre direction than along the base S^2 directions.

In Figure 3.4, we study the behaviour of the energy density as a function of the radial coordinate

R in the $R \to \infty$ region. As one can see from the figure, the energy density decays like $1/R^7$ near infinity, and the other components of the stress tensor exhibit the same asymptotic behaviour. Unsurprisingly, this can be understood in a similar manner to the 4D case [81]. The energy density stored in the field around a point particle in 5-spacetime dimensions behaves like $\propto N_c^{\frac{5}{2}} R_g/R^7$, where R_g is the gravitational radius of the point particle. The pressures are comparable in magnitude to the energy density and also decay like $1/R^7$. Note that these results agree with the braneworld calculation presented in the previous section.

Finally, we consider the rotation of the dual plasma. To do so, we write the unique timelike eigenvector of the stress energy tensor of the dual CFT as

$$T = \frac{\partial}{\partial t} + \Omega(R) \frac{\partial}{\partial \psi} . \tag{3.53}$$

This allows us to identify $\Omega(R)$ as the angular velocity of the plasma with respect to a static observer at infinity in the boundary directions. In Figure 3.5 we plot this quantity for different values of the angular momentum parameter a. As one might expect, $\Omega(R) \to \Omega_H$ as $R \to R_H$. That is, at the horizon the plasma is co-rotating with the black hole. This is an example of frame dragging, even though at the boundary gravity is non-dynamical. This result had to hold, since the boundary black hole is rotating rigidly with respect to infinity; there is no flux of CFT plasma through the horizon, so the plasma must be co-rotating with the black hole. In the asymptotic region $R \to \infty$, we find that the angular velocity of the plasma decays like $1/R^2$, so that at infinity the plasma is at rest and the timelike eigenvector coincides with the asymptotic timelike Killing vector.

3.3.4 Negative energy density

Figure 3.3 in the previous section clearly showed that the energy density of the plasma is negative. Therefore (3.52) violates the weak and strong energy conditions, which are what is usually considered physical in classical general relativity. This is also the case for the stress-energy tensor of $\mathcal{N} = 4$ SYM in the background of a 4D Schwarzschild black hole in the Unruh vacuum [48]. However, the energy density in the Hartle–Hawking state was found to be positive in [62]. While a negative energy density may be a sign of a pathology and/or instability in classical solutions, it

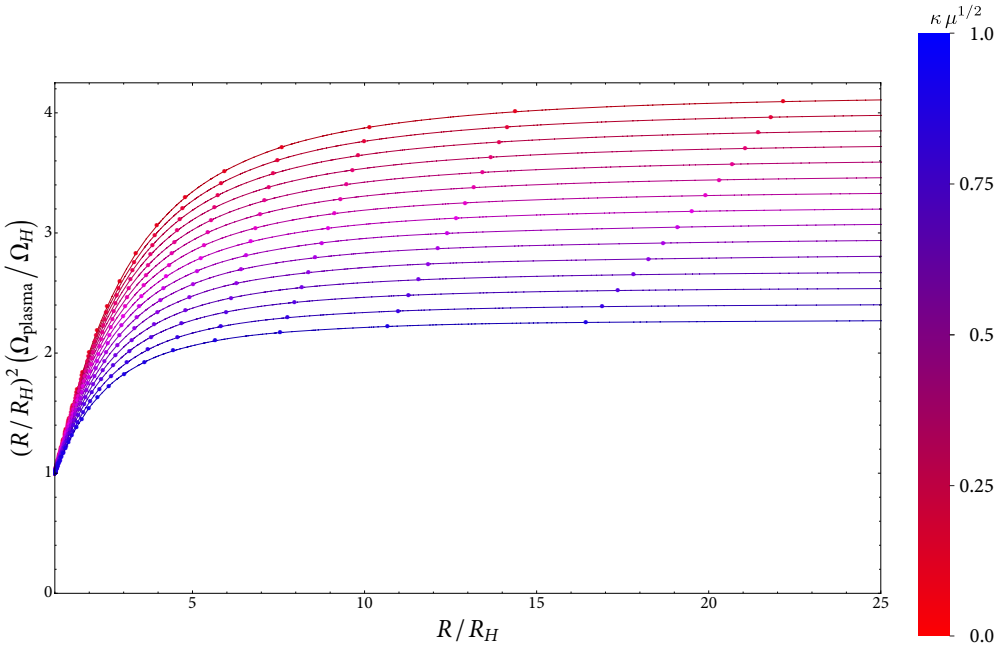

Figure 3.5: Rescaled angular velocity of the plasma as a function of the radial coordinate R for different values of the rotation parameter. Near infinity the angular velocity of the plasma decays as $\sim 1/R^2$, while at the horizon it reduces to the angular velocity of the black hole.

is not necessarily problematic in our present setting, as we are considering a *quantum* field theory.

References [48, 81] provide a heuristic physical interpretation of the stress tensor (3.52) as a 'halo' of plasma in equilibrium with the boundary black hole. According to this picture, the strong attractive self-interactions of the CFT are balanced by thermal radiation pressure from the black hole. Another instance in which a non-trivial background results in a QFT stress tensor with similar qualitative properties is the celebrated Casimir effect [82]. In that setting, the vacuum energy of the electromagnetic field between two perfectly conducting plates is negative, giving rise to an attractive force between the two plates. Reference [68] also noted a large contribution of the Casimir effect in the stress tensor of a conformally coupled weakly interacting scalar field in the background of the BTZ black hole. We may therefore be able to attribute our negative energy density in (3.52) to the attractive nature of the self-interactions of the CFT due to the Casimir effect.

In the calculations that follow, we can concentrate only on the t-r part of the geometry, as this gives rise to the non-trivial causal structure of the spacetime. Under our symmetry assumptions,

the sphere directions are merely spectators. We begin by analysing the t-t and r-r components of the stress tensor through a simple 2D model. For concreteness we consider $\langle T^i{}_j \rangle$ for a generic quantum field theory in the 2D space

$$\mathrm{d}s^2 = -\left(1 - \frac{2\,M}{r}\right)\mathrm{d}t^2 + \frac{\mathrm{d}r^2}{1 - \frac{2\,M}{r}}\,. \tag{3.54}$$

The advantage of working in 2D is that the stress tensor is completely determined by symmetries and conservation in terms of purely geometrical data [83, 84]. Indeed, assuming that the stress tensor is time-symmetric, one finds that the conservation equation reduces to

$$\partial_r \langle T^r{}_r \rangle = \frac{M}{r^2}\frac{1}{1 - \frac{2\,M}{r}}\big(\langle T^t{}_t \rangle - \langle T^r{}_r \rangle\big) \quad \Rightarrow \quad \partial_r \left[\left(1 - \frac{2\,M}{r}\right)\langle T^r{}_r \rangle\right] = \frac{M}{r^2}\,\langle T^i{}_i \rangle\,. \tag{3.55}$$

This equation can be readily integrated to get

$$\langle T^i{}_j \rangle = \mathrm{diag}\left(-\frac{H_2(r)}{1 - \frac{2\,M}{r}} + \langle T^i{}_i \rangle,\, \frac{H_2(r)}{1 - \frac{2\,M}{r}}\right)\,, \tag{3.56}$$

where the first entry denotes the t-t component of the stress tensor, and

$$H_2(r) := M\int_{2M}^{r} \mathrm{d}r'\,\frac{\langle T^i{}_i \rangle(r')}{r'^2}\,. \tag{3.57}$$

The full stress tensor is now completely determined by its trace, which is fixed by the conformal anomaly. For the 2D geometry (3.54), this is given by

$$\langle T^i{}_i \rangle = \alpha\, R = \frac{4\,\alpha\, M}{r^3}\,, \tag{3.58}$$

where α is a constant that depends on the spin of the field being considered. For a free massless scalar field, we have $\alpha = \frac{1}{24\,\pi}$ [83, 84]. The stress tensor in this geometry is therefore

$$\langle T^i{}_j \rangle = \alpha\,\mathrm{diag}\left(\frac{4\,M}{r^3} - \frac{1}{16\,M^2}\big(1 + \tfrac{2M}{r}\big)\big(1 + \tfrac{4M^2}{r^2}\big),\, \frac{1}{16\,M^2}\big(1 + \tfrac{2M}{r}\big)\big(1 + \tfrac{4M^2}{r^2}\big)\right)\,. \tag{3.59}$$

In the near horizon region, (3.59) has the same qualitative features as our 5D stress tensor. In particular, it is static and regular on both the future and past event horizons, and the energy density near the horizon is *negative*. Since the origin of (3.59) is clearly the near-horizon spacetime

curvature itself, we can conclude that it is a vacuum polarisation effect. However, this 2D model does not capture the physics in the asymptotic region near infinity. For large r (3.59) reduces to the stress tensor of radiation at the Hawking temperature set by the horizon in (3.54), and this is not what we observe in our numerical results.

We now add additional dimensions and consider the full 4D Schwarzschild geometry. As we are only interested in CFTs in this background, we assume that the stress tensor is static, spherically symmetric and traceless. With these assumptions, conservation of the stress tensor fully determines its components in terms of one undetermined function [84, 48]:

$$\begin{aligned} \langle T^{i}{}_{j} \rangle &= \text{diag}\left(-\frac{H_4(r)}{r^2(1-\frac{2M}{r})} - 2\,\Theta(r),\ \frac{H_4(r)}{r^2(1-\frac{2M}{r})},\ \Theta(r),\ \Theta(r) \right) , \\ H_4(r) &= 2\int_{2M}^{r} \mathrm{d}r'(r'-3\,M)\Theta(r') \,, \end{aligned} \tag{3.60}$$

where the first two entries in (3.60) correspond to the t-t and r-r components and $\langle T^{\theta}{}_{\theta} \rangle = \langle T^{\phi}{}_{\phi} \rangle = \Theta(r)$ by spherical symmetry. Note that the extra dimensions introduce new pressures along the angular directions, and the extra powers of r in the denominators of (3.60) can alter the behaviour of $\langle T^{i}{}_{j} \rangle$ near infinity.

The actual form of the function $\Theta(r)$ in (3.60) depends on the details of the CFT under consideration. However, one can constrain it on general grounds without assuming strong coupling. First, note that for (3.60) to be finite at the horizon ($r = 2\,M$), we only require that $\Theta|_{r=2M}$ be finite. In addition, we saw above that the 2D model captures the vacuum polarisation effect near the horizon, and therefore we require that $\langle T^{t}{}_{t} \rangle > 0$ for $r \approx 2\,M$. This implies that $\Theta < 0$ in this region, and so we conclude that the polarisation of the vacuum near the horizon will inevitably make the pressures along the angular directions negative. Also note that from (3.60) we can conclude that the following relations must hold at the horizon: $\langle T^{t}{}_{t} \rangle|_{r=2M} = \langle T^{r}{}_{r} \rangle|_{r=2M} = -\Theta|_{r=2M}$. A close examination of the data of [48] shows that this is true for $\mathcal{N} = 4$ SYM in the background of Schwarzschild, the deviations being smaller than 1%. At infinity we want $\langle T^{i}{}_{j} \rangle \to 0$, and for this we require $\Theta \sim 1/r^3$ or faster. Then (3.60) implies that all the components of the stress tensor have the same behaviour near infinity. Note that Θ has to decay sufficiently fast if we require that $\langle T^{t}{}_{t} \rangle$ and $\langle T^{r}{}_{r} \rangle$ do not change sign at a sufficiently large r. One can infer the fall off of Θ from the linearised calculation of [72] and one finds that the pressures along the sphere directions are

negative and the fall off is $\sim 1/r^5$.

Finally, we note that negative energy densities can be bound by quantum energy inequalities (see [85] for a review with references). These are the remnants of the classical energy conditions satisfied by quantum field theories, and they quantify the extent to which a quantum field can violate these classical energy conditions. Even though most of the results available in the literature apply only to Minkowski spaces, one might expect similar results to apply on curved spacetimes over regions which are small compared the curvature length scales. Ref. [86] shows that this intuition is indeed correct. Our results, obtained using AdS/CFT and valid at strong coupling, show that in this regime the energy density also remains bounded below.

3.4 Conclusions

In this chapter, we constructed the gravitational duals of 5D CFTs in rotating black hole backgrounds. For simplicity, we considered only the case in which the boundary black hole is the 5D Myers–Perry solution with equal angular momenta. Far from the horizon, the spacetime reduces to the Poincaré horizon of AdS, and hence the CFT is in the Unruh state. While our solutions should capture the general qualitative physics of the Unruh vacuum in rotating black hole backgrounds (at large N_c and strong coupling), it would still be interesting to consider the gravitational dual of $\mathcal{N} = 4$ SYM on the background of Kerr. From a computational perspective, the problem reduces to solving elliptic PDEs in 3 variables, which is considerably more expensive than the work we present here. The construction of gravitational duals to CFTs on rotating black hole backgrounds becomes even more interesting if one considers an IR horizon with a finite temperature. As [62] pointed out, this requires that the horizon of the black hole be non-Killing. We noted that a generic feature of the Unruh vacuum for strongly coupled CFTs in black hole backgrounds is that the energy density is everywhere negative in the domain of outer communications. We argued that this can be understood in terms of the polarisation of the vacuum due to the curvature of the background. Furthermore, our stress tensor (3.51) is bounded from below, in accordance with generic local quantum energy inequalities [85]. The fact that the energy density is negative need not signal that our solutions are unstable, but we did not perform a detailed analysis of their stability under gravitational perturbations.

It is interesting to compare the results of this chapter (and of [48]) with black funnels [62]. The latter represent the gravitational dual of the Hartle–Hawking state for $\mathcal{N} = 4$ SYM on the background of Schwarzschild. In the black funnels case the energy density is positive *everywhere* (see Figure 4 in [62]), and not only at infinity, where the stress tensor reduces to that of pure radiation at the Hawking temperature set by the IR horizon. Therefore, the stress tensor does not seem to have a component corresponding to the vacuum polarisation. This can be contrasted with the free field theory results [87], which indicate that near the horizon it is the vacuum polarisation that dominates the stress tensor, as in this paper. However, [59] conjectured that there should exist another family of solutions dual to the Hartle–Hawking state, namely the black droplets. The latter should arise as a continuous deformation of the solutions of [48], where the continuous parameter is the temperature of the IR horizon. Because in [48] (and this work) the non-vanishing stress tensor is due to vacuum polarisation effects, the stress tensor of the black droplets should be dominated by vacuum polarisation effects near the horizon, at least when the temperature of the IR black hole is sufficiently small relative to the size of the boundary black hole. In the asymptotic region, the stress tensor should reduce to pure radiation, as in the black funnels case. Therefore, for black droplets the $\langle T^t_{\ t} \rangle$ component of the stress tensor should change sign, being positive near the horizon and negative in the near infinity region.

CHAPTER 4

Localised Plasma Balls

The work presented in this chapter was published as [88]. The paper was co-authored with Pau Figueras. Both authors wrote separate code for this project, which allowed us to cross-check results. The results should be regarded as a product of collaborative work.

4.1 Introduction

AdS/CFT can be applied to the study of the confinement/deconfinement phase transition in strongly coupled gauge theories. This was originally proposed in [41]. Conformal symmetry implies that a CFT in Minkowski space cannot have a phase transition at any finite temperature. Since any non-zero temperature can be scaled to any other value, all non-zero temperature states must be equivalent. From the bulk perspective, this translates into the fact that the planar Schwarzschild–AdS solution is the dominant phase for all temperatures. When conformal symmetry is broken by putting the CFT on a sphere, one has a phase transition which has been identified as a confinement/deconfinement phase transition. The bulk dual of this phase transition corresponds to the celebrated Hawking–Page phase transition between thermal AdS and global Schwarzschild–AdS [89]. Strictly speaking, this transition is only present in the $N_c \to \infty$ limit. Ref. [90] provides a much more extensive review on this and related topics.

In this chapter, we will study the confinement/deconfinement transition of CFTs on a flat space with one spatial direction compactified onto a circle. Fermions are subject to anti-periodic (Scherk–Schwarz) boundary conditions on this circle. The compact circle breaks conformal invariance,

therefore allowing the CFT to have a non-trivial confinement/deconfinement phase transition as the temperature changes. From the bulk perspective, the low temperature confining vacuum corresponds to the *AdS soliton* geometry [41, 91], while the high temperature deconfined phase corresponds to the planar Schwarzschild–AdS geometry. Further details of these geometries are reviewed in §4.2.1.

At the deconfinement temperature, the period of the Euclidean time circle and that of the Scherk–Schwarz circle coincide, while the pressure of the plasma vanishes. This suggests the existence of a new solution to the Einstein equations in the bulk which interpolates between the AdS soliton geometry and the planar Schwarzschild–AdS black hole. This solution would be dual to a domain wall configuration of the CFT. Indeed, such a solution has been numerically constructed in [42]. The authors of this work further conjectured that finite size black holes, localised at the IR bottom of the AdS-soliton background, should also exist. From the dual CFT perspective, these black holes would correspond to localised balls of deconfined plasma surrounded by the confining vacuum. In fact, [42] argued that such finite size plasma balls should generically exist in any large N_c gauge theory that exhibits a first order confinement/deconfinement phase transition. In the semiclassical approximation, these plasma balls should be stable, as in the full quantum theory they can only evaporate via the emission of colour singlet glueball states; a process which is dual to Hawking evaporation of the bulk black hole. However, out of the $O(N_c^2)$ degrees of freedom available in the theory, only $O(1)$ of them correspond to the colour singlet states that can be emitted into the confining vacuum, thus suppressing the evaporation process in the large N_c limit [42].

The work in [92] took the first steps towards a quantitative understanding of plasma balls in confining backgrounds. The authors considered a particular form of the four-dimensional AdS C-metric [93], which corresponds to a black hole moving in an accelerated trajectory inside AdS, but with the geometry cut off in the IR. This construction can be viewed as the complement of the one in [94], which constructed localized black holes on the brane in Randall-Sundrum II braneworld models. The solutions of [92] confirmed some of the predictions made in [42]. In particular, the bulk horizon corresponding to large plasma balls does indeed have a pancake-like shape that extends along the IR bottom of the geometry. The properties of the plasma at the interior of the ball are also those of the deconfined state at the same temperature. However, the solutions of [92] also exhibited certain peculiarities due to the boundary conditions at both the IR and UV ends of

the geometry. For instance, the boundary metric is not flat and asymptotes to a cone with a certain excess angle.

In light of these results, it is desirable to investigate the physical properties of these plasma balls in a cleaner setting, such as the framework of [42]. In this chapter, we employ numerical methods to construct localized black holes at the IR bottom of the AdS soliton geometry, which are dual to the sought plasma balls. For simplicity, we consider static and spherically symmetric configurations from the point of view of the non-compact directions along the boundary. This reduces the problem to solving non-linear PDEs in two variables. Here, we limit our construction to 5 and 6 bulk dimensions; in either case, the qualitative physics of the plasma balls is the same. Having found solutions, we proceed to study some of their physical properties. First, we study the shape of the horizon for plasma balls of different sizes. We find that small plasma balls are described by approximately spherical black holes in the bulk, while large plasma balls are dual to pancake-like black holes which extend along the IR bottom of the geometry. Away from the edges of the black hole, the geometry is well-described by a homogeneous black brane at the same temperature. By extracting the stress of the dual field theory, we can also assess the behaviour as the deconfinement temperature is approached. Under these conditions, the energy density at the centre of the ball approaches the energy density of a homogeneous black brane at the same temperature. For large plasma balls, the region in the vicinity of their edge reduces to a good approximation to the domain wall solution of [42]. In particular, we find that the tension is positive. We should also note that in this chapter we are only considering solutions in the universal gravity sector. In the presence of a compact internal space, we should expect that plasma balls which are much smaller than the length scale of the internal space will be subject to the Gregory–Laflamme instability.

The chapter is organized as follows. In §4.2 we explain our set up. §4.2.1 reviews the confinement/deconfinement phase transition in relevant backgrounds, while §4.2.2 gives details of our numerical construction of black holes which are dual to plasma balls. §4.3 is devoted to analysing the physical properties of plasma balls. In §4.3.1 we study the horizon geometry of the black holes and in §4.3.2 we analyse temporal Wilson loops as a probe of confinement. Finally, §4.3.3 presents our results for the stress tensor of the dual field theory.

4.2 Numerical setup

This section explains our numerical construction of localized plasmaballs. In §4.2.1 we review the necessary features of both the AdS-soliton metric and the black brane. In §4.2.2 and §4.2.2 we then explain the details of the actual numerical construction.

4.2.1 Preliminaries

Here, we briefly review the basics of thermal field theory in Scherk–Schwarz–AdS (SS–AdS), with anti-periodic boundary conditions for the fermions on the SS circle. For further details, see the recent review [90] and references therein.

We are interested in solutions to the Einstein vacuum equations in $D+1$ dimensions with SS–AdS boundary conditions. In other words, we require that the boundary geometry is conformal to $\mathbb{R}^{1,D-2} \times S^1_{\rm SS}$. We will take the fermions in the CFT to be anti-periodic on the circle $S^1_{\rm SS}$, which is usually referred to as the Scherk–Schwarz (SS) circle. The anti-periodic boundary condition is required for the existence of a spin structure on bulk manifolds whose SS circle shrinks to zero size, such as the AdS-soliton solution [41]. Since we wish to use AdS/CFT to study thermal field theories, we will work with the Euclidean section $t = -\mathrm{i}\,\tau$. The Euclidean time τ is periodic with $\tau \sim \tau + \beta$, where the period β is the inverse temperature. Therefore, in this chapter, we consider $(D+1)$-dimensional asymptotically locally AdS spaces whose boundary metric is $\mathbb{R}^{D-2} \times S^1_\beta \times S^1_{\rm SS}$. Since we will only address static configurations here, working in the Euclidean section does not make any practical difference at the level of the numerical construction.

Within this class of geometries, it has been conjectured that the so-called AdS soliton [41, 91] is the actual ground state [91]:

$$\mathrm{d}s^2_{\rm soliton} = \frac{\ell^2}{z^2}\left(\mathrm{d}\tau^2 + \frac{1}{f(z)}\,\mathrm{d}z^2 + \mathrm{d}\mathbf{x}^2_{D-2} + f(z)\,\mathrm{d}\theta^2\right), \qquad f(z) = 1 - \left(\frac{z}{z_0}\right)^D, \tag{4.1}$$

where $d\mathbf{x}^2_{D-2} = \sum_{i=1}^{D-2} dx^i\, dx^i$ is the flat metric on $\mathbb{R}^{D-2}$. The space-time (4.1) is smooth and

complete if and only if the period of the SS circle is given by

$$\Delta\theta = \frac{4\,\pi\,z_0}{D} \equiv L\,. \tag{4.2}$$

Near $z = z_0$ the metric (4.1) approaches the flat metric on $\mathbb{R}^{D-2} \times S^1_\beta \times \mathbb{R}^2$ and therefore the global topology of the space-time is $S^1_\beta \times \mathbb{R}^D$. This shrinking of the SS circle in the bulk smoothly cuts off the geometry in the IR, leading to a mass gap and confinement [41]. Note that in (4.1) the period of the Euclidean time circle can be arbitrary. Thus, from the point of view of the canonical ensemble, the AdS-soliton phase (4.1) exists at all temperatures.

By exchanging the thermal and the SS circles (i.e. $\tau \leftrightarrow \theta$) in (4.1), we can immediately write down another solution to the Einstein equations obeying the same boundary conditions:

$$\mathrm{d}s^2_{brane} = \frac{\ell^2}{z^2}\left(f_\beta(z)\,\mathrm{d}\tau^2 + \frac{1}{f_\beta(z)}\,\mathrm{d}z^2 + \mathrm{d}\mathbf{x}^2 + \mathrm{d}\theta^2\right), \qquad f_\beta(z) = 1 - \left(\frac{z}{z_h}\right)^D, \tag{4.3}$$

This is the well-known planar Schwarzschild–AdS black hole. In this solution, the Euclidean time circle is contractible in the bulk, and the regularity of the geometry (4.3) at $z = z_h$ fixes the temperature as a function of the IR cutoff:

$$\beta = \frac{4\,\pi\,z_h}{D}\,. \tag{4.4}$$

Note that in this case the SS circle is non-contractible and hence the period of θ can be arbitrary.

There is yet another classical solution to the Einstein equations obeying these boundary conditions, namely thermal AdS, which is simply pure Poincaré–AdS space-time with both τ and θ suitably identified. However, the free energy of this saddle point is always greater than the free energies of the other saddles (the AdS-soliton and the planar Schwarzschild-AdS black hole). Therefore, thermal AdS never dominates the canonical ensemble and thus we ignore it from now on.

In this chapter we are interested in studying the finite temperature phases in SS–AdS. By comparing thermodynamic quantities, in particular the free energy, one can determine which is the dominant phase at a given temperature. It turns out that at temperatures $T < \frac{1}{L}$, the AdS-soliton has the lowest free energy and is therefore the dominant phase. By contrast, at temperatures $T > \frac{1}{L}$ the planar black hole dominates the thermal ensemble, and at $T = T_d = \frac{1}{L}$ there is a first order

confinement/deconfinement phase transition. Given the symmetry between the τ and θ circles in the AdS-soliton and planar black hole geometries, it is unsurprising that the phase transition occurs precisely when the sizes of these two circles become equal. In fact, at $T = T_d$, (4.1) and (4.3) are symmetric under exchange of the τ and θ circles. This led the authors of [42] to construct numerically a domain wall solution at temperature $T = T_d$ that interpolates between the confined and deconfined regions. The wall has thickness $\sim L$, since at $T = T_d$ the thermal scale and the scale of the mass gap coincide. One can also measure the tension of the wall μ_{d-1} in the field theory, which turns out to be positive. We will return to this point when we analyse the physics of plasma balls in Section 4.3.3.

The existence of this domain wall solution led the authors of [42] to conjecture the existence of finite size black holes localized in the IR region of the geometry. From the perspective of the dual CFT, these black holes would correspond to bubbles of deconfined plasma within the confining vacuum. More precisely, [42] conjectured that generic confining backgrounds should host a one-parameter family of black holes, labeled by a mass, which are spherically symmetric in the $p = D - 2$ non-compact spatial dimensions, and which satisfy the following properties:

1. The radius of these black holes in p dimensions scales with the mass like $m^{\frac{1}{p}}$.
2. Away from the edge in the p non-compact dimensions, the geometry of these black holes approximate the black brane at $T = T_d$.).
3. In the vicinity of the edge, these black holes reduce to a domain wall, which interpolates between a black brane at T_d and the AdS soliton.

In this chapter we will provide strong evidence that such black holes do exist by numerically constructing them. Moreover, by extracting some of their physical properties we will be able to test certain conjectures put forward in [42].

Before moving on to describe the details of our numerical construction we will outline some conventions. Unless otherwise stated, we will set the AdS radius ℓ to one. In order to study the thermal phases in SS–AdS we first recall that by conformal invariance only the dimensionless ratio

β/L is physical. It is therefore convenient to fix the scale by setting

$$z_0 = 1 \quad \Rightarrow \quad L = \frac{4\pi}{D}\,, \tag{4.5}$$

and then vary the inverse temperature β. As such, the deconfinement temperature in our units is $T_d = \frac{D}{4\pi}$.

4.2.2 Numerical details

To construct these black holes numerically, we use the Einstein–DeTurck method as described in Section 2.2. Ref. [48] showed that all solutions to (2.8) are necessarily Einstein for static space-times that are asymptotically flat, AdS or Kaluza–Klein, and whose boundary conditions are compatible with ξ^a vanishing at the boundaries $\partial\mathcal{M}$ of the manifold. For the problem we consider in this chapter, these conditions are met, and therefore solving (2.8) is equivalent to solving the Einstein equations.

We seek finite size black holes which are asymptotically SS–AdS, and which are localized in the IR of the geometry. This can be thought of as finite energy (as opposed to finite energy density) excitations about the AdS soliton background. For simplicity, we consider only static black holes.

The AdS soliton is the lowest energy state in the class of metrics with boundary geometry $\mathbb{R}^{D-2} \times S^1_\beta \times S^1_{\rm SS}$ [91]. Since the global topology of the spatial sections of the background is $\mathbb{R}^D$, one expects that localized black holes in the bulk should exist. These would be static black holes that are asymptotically SS–AdS, and have a horizon with S^{D-1} spatial topology. Indeed, in the limit in which the black hole is much smaller than the AdS radius and the SS circle radius at infinity, one would expect such solutions to be well-approximated by the standard asymptotically flat Schwarzschild black hole. In §4.3.1, we will show that this intuition is indeed correct. From the dual CFT perspective, these black holes correspond to states that are localized on the $\mathbb{R}^{D-2}$ part of the boundary geometry, and not on the $S^1_{\rm SS}$, as this is a contractible circle in the bulk. For simplicity, we restrict ourselves to bulk space-times which are rotationally symmetric from the point of view of the boundary, in the sense that they preserve the $SO(D-2) \times U(1)_{\rm SS}$ symmetry of the spatial $\mathbb{R}^{D-2} \times S^1_{\rm SS}$ flat boundary metric. From the perspective of the CFT, these black

holes should correspond to the plasma balls considered in [42]; i.e., rotationally symmetric balls of deconfined plasma sitting in the confining vacuum.

Metric ansatz and boundary conditions

It is very difficult to construct a single coordinate system that is adapted to the entire geometry under our symmetry assumptions. We work around this difficulty by instead using two coordinate charts: one adapted to the asymptotic region far from the black hole and the other adapted to the near-horizon region. This approach was first used successfully in the numerical construction of 5D localized Kaluza–Klein black holes [45].

In the region far from the bulk black hole, the space-time should approach the AdS soliton metric (4.1). Since we are interested in preserving the spatial $SO(d-2)\times U(1)_{\rm SS}$ symmetry of the boundary metric, we rewrite (4.1) to make these symmetries manifest as

$$\mathrm{d}s^2_{\rm soliton} = \frac{1}{z^2}\left(\mathrm{d}\tau^2 + \frac{1}{f(z)}\,\mathrm{d}z^2 + \mathrm{d}\rho^2 + \rho^2\,\mathrm{d}\Omega^2_{(D-3)} + f(z)\,\mathrm{d}\theta^2\right), \tag{4.6}$$

where $f(z)$ has already been defined. For computational purposes, we find it convenient to introduce new compact coordinates (x, y),

$$z = 1 - y^2\,, \qquad \rho = \frac{k_x\, x}{1-x^2}\,, \tag{4.7}$$

where k_x is a freely adjustable parameter that allows us to stretch the ρ coordinate. With these co-ordinates, we can bring the asymptotic region $\rho\to\infty$ in to a finite coordinate distance $x = 1$. In terms of these new coordinates, the AdS soliton metric becomes,

$$\mathrm{d}s^2_{\rm soliton} = \frac{1}{(1-y^2)}\left(\mathrm{d}\tau^2 + \frac{4\,z_0^2}{\bar f(y)}\,\mathrm{d}y^2 + \frac{k_x^2(1+x^2)^2}{(1-x^2)^4}\,\mathrm{d}x^2 + \frac{k_x^2\,x^2}{(1-x^2)^2}\,\mathrm{d}\Omega^2_{(D-3)} + \bar f(y)\,\mathrm{d}\theta^2\right), \tag{4.8}$$

where we have defined $\bar f(y)$ via $f(y) = y^2\,\bar f(y)$. We can now write down the ansatz for the

metric in the outer region:

$$\mathrm{d}s^2_{\text{outer}} = \frac{1}{(1-y^2)^2}\Bigg(T\,\mathrm{d}\tau^2 + \frac{4\,y^2\,\bar{f}(y)\,S}{D^2}\,\mathrm{d}\theta^2 + \frac{k_x^2\,x^2\,R}{(1-x^2)^2}\,\mathrm{d}\Omega^2_{(D-3)} + \frac{k_x^2(1+x^2)^2\,A}{(1-x^2)^4}\,\mathrm{d}x^2 + \frac{4\,B}{\bar{f}(y)}\,\mathrm{d}y^2 - \frac{2\,k_x\,(1+x^2)\,F}{(1-x^2)^2(1-y^2)}\,\mathrm{d}x\,\mathrm{d}y\Bigg)\,. \tag{4.9}$$

Here, the functions $\{T,\ R,\ S,\ A,\ B,\ F\}$ are our unknowns and they depend on both x and y. In these coordinates, $x = 0$ and $y = 0$ correspond to the fixed point sets of the $SO(D-2)$ symmetry and the $U(1)_{\text{SS}}$ symmetry respectively. Regularity at these axes requires that all functions satisfy a Neumann condition there, except for F, which must satisfy a Dirichlet boundary condition. To avoid conical singularities at $x = 0$ we must also impose $A = R$; similarly, the avoidance of conical singularities at $y = 0$ requires that $B = S$ there. The boundary $x = 1$ is an asymptotic end, at which the space-time should approach the AdS soliton metric (4.8). Therefore, we impose Dirichlet boundary conditions there: $T = R = S = A = B = 1$, $F = 0$. Similarly, $y = 1$ is the boundary of AdS and we impose Dirichlet boundary conditions there as before. Recall that by varying the norm of the Euclidean time circle at the boundary we can change the temperature of the bulk black hole. We will make use of this in our construction to efficiently explore the branch of solutions. Finally, note that (4.9) is not adapted to describe the horizon in the sense that in these (x, y) coordinates the horizon does not lie at a constant value of either of them. Therefore, in order to avoid unnecessary difficulties when dealing with horizons in non-adapted coordinates, we ensure that the domain covered by (4.9) does not contain horizons. We will return to this point in Section 4.2.2, when we describe our computational domain.

In the region near the horizon, the space-time should approach the near-horizon region of a topologically spherical black hole. We expect that for small plasma balls, the corresponding bulk black holes should approximate the asymptotically flat, static and round black holes. The geometry should therefore be close to spherical symmetry in the full $(D+1)$-dimensional sense. For larger plasma balls, the presence of a non-zero cosmological constant and the non-trivial topology of the IR bottom of the geometry should become significant. The black holes should therefore be highly deformed from spherical symmetry. It is clearly desirable to have an ansatz which can be adapted

to suit each of these two regimes. To achieve this, we write the near-horizon ansatz as

$$\begin{aligned} \mathrm{d}s^2_{\text{inner}} = \frac{1}{(1-r^2)^2}\Big[& r^2\, g(r)\, T'\, \mathrm{d}\tau^2 + r_0^2\Big((1-a^2)^2\, R'\, \mathrm{d}\Omega^2_{(D-3)} + a^2(2-a^2)\, S'\, \mathrm{d}\theta^2\Big) \\ & + \frac{4\, r_0^2\, A'}{g(r)}\, \mathrm{d}r^2 + \frac{4\, r_0^2\, B'}{2-a^2}\, \mathrm{d}a^2 + 2\, r\, F'\, \mathrm{d}r\, \mathrm{d}a\Big]. \end{aligned} \tag{4.10}$$

where $g(r)$ is a freely specifiable function. The functions $\{T',\ R',\ S',\ A',\ B',\ F'\}$ are our unknowns. The radial coordinate r in (4.10) ranges from $r = 0$ (the horizon) up to some $r = r_{\text{outer}}$ that we can in principle freely choose. In practice, r_{out} can neither be too large nor too small, and we find that $r_{\text{out}} \simeq 0.6$ works well. In our calculations, we choose $g(r)$ so that the near-horizon geometry (4.10) coincides with the asymptotically AdS Schwarzschild black hole when $T' = R' = S' = A' = B' = 1$ and $F' = 0$. Regularity at the horizon requires that all functions satisfy the Neumann condition there. In addition, we must impose $T' = A'$ to ensure that the space-time metric and the reference metric have the same surface gravity. Here, r_0 is a dimensionless parameter related to the surface gravity by

$$\kappa^2 = \frac{g(0)^2}{4\, r_0^2}\,. \tag{4.11}$$

a is an angular coordinate, whose range is $0 \leq a \leq 1$. At $a = 0$ the SS circle shrinks to zero and regularity requires that all functions satisfy the Neumann condition there except for F', which must vanish. To avoid conical singularities we also impose $B' = S'$. Similarly, at $a = 1$ it is the $(d-3)$-sphere that shrinks to zero size; regularity requires that all functions satisfy a Neumann boundary condition there and that $F' = 0$. Absence of conical singularities further imposes $B' = R'$ at $a = 1$.

Computational domain and reference metric

In our construction, the far region coordinates (x, y) and the near region coordinates (r, a) are simply related by

$$x = r\,(1-a^2)\,, \qquad y = r\, a\, \sqrt{2-a^2}\,. \tag{4.12}$$

Note that by construction the range of the far region coordinates is $0 \leq \{x\,, y\} \leq 1$. This condition limits the range of r in the inner region patch. In terms of the far region (x, y) coordinates, the

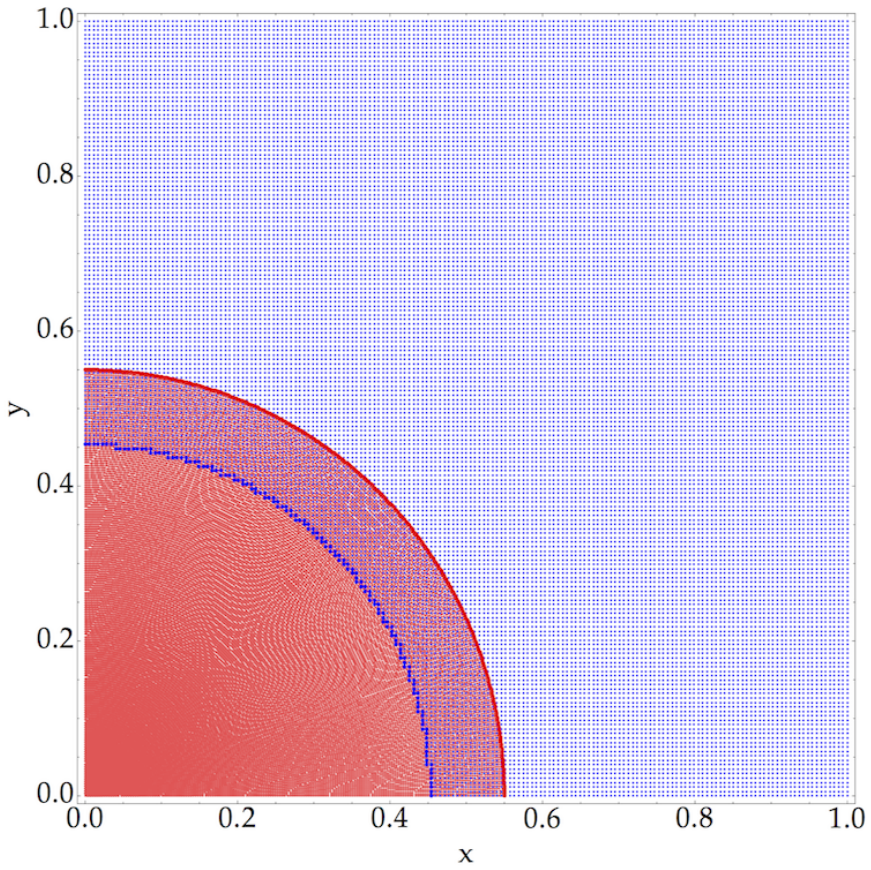

Figure 4.1: Computational domain. The blue dots correspond to grid points in the far region patch, while the red dots are grid points in the near region patch, expressed in the far region coordinates. The larger dots indicate boundary points whose values are obtained by interpolation of the metric functions in the other patch.

horizon ($r = 0$) appears to be point-like, but this is only a coordinate artefact.

We solve (2.8) numerically using sixth-order finite differences and Newton's method. One advantage of using finite differences is that we can cover our domain by two simple patches. Figure 4.1 displays our computational domain and shows the patches: the far patch (in blue) and the near patch (in red). As this figure indicates, the two coordinate charts overlap in a certain region; by appealing to the uniqueness of the solution in the whole domain, we can use this overlapping region to transfer information between the values of the metric functions in the two patches. Although in principle the size of the overlapping region is arbitrary, in practice we found that a relatively small overlap between the patches, as in Figure 4.1, worked best. Function values on the *interpolation boundary* of the overlapping region (the points indicated by larger dots in Figure 4.1) are obtained

by interpolating values taken from the other patch. More precisely, at each point p on the interpolation boundary, we choose in the other patch a 6x6 square of points whose centre is closest to p. We can uniquely fit a local polynomial of fifth degree to the six chosen points in each grid direction from which we can read off approximate values for the functions at p. Although this high order interpolation is still likely to introduce a local error with a lower order convergence than that of our differentiation scheme, we have not observed a loss of convergence order within the range of resolutions employed. The resolutions used appear to be sufficient for extracting the physical quantities of interest.

In order to solve (2.8), we need to provide a global reference metric in the space-time manifold as part of the gauge fixing procedure. Following [45], we achieve this by a simple interpolation between the near-horizon geometry and the AdS soliton metric (4.8):

$$\bar{g}_{\mu\nu}\,\mathrm{d}x^\mu\,\mathrm{d}x^\nu = \left[\left(1 - I(r; d_{\min}, d_{\max})\right)\bar{g}^{\text{inner}}_{\mu\nu} + I(r; d_{\min}, d_{\max}))\,\bar{g}^{\text{outer}}_{\mu\nu}\right]\mathrm{d}x^\mu\,\mathrm{d}x^\nu\,, \tag{4.13}$$

where we have used the interpolation function

$$I(r; d_{\min}, d_{\max}) = \begin{cases} 0 & , r \le d_{\min} \\ \frac{1}{2} - \frac{1}{2}\tanh\left[\cot\left(\pi\left(1 + \frac{d_{\max}-r}{d_{\max}-d_{\min}}\right)\right)\right] & , d_{\min} < r < d_{\max} \\ 1 & , d_{\max} \le r \end{cases} \tag{4.14}$$

and $d_{\min}, d_{\max}$ are freely adjustable parameters. These parameters control the size of the overlapping region between the near-horizon geometry and the AdS-soliton metric. In our calculations we typically used $d_{\min} = 0.2$ and $d_{\max} = 0.9$. Note that (4.14) is a smooth compactly supported function; this is useful to ensure that the reference metric (4.13) satisfies our boundary conditions in all regions. In (4.13), $\bar{g}^{\text{outer}}$ denotes the AdS soliton metric, which can be obtained from (4.9) by setting $T = R = S = A = B = 1$, $F = 0$; similarly, $\bar{g}^{\text{inner}}$ denotes the near-horizon black hole metric and is obtained from (4.10) by setting $T' = R' = S' = A' = B' = 1$, $F' = 0$.

As conjectured in [42], and as we confirm here, there exists a unique branch of plasma balls labelled by temperature. As such, we can explore the branch of solutions by varying the parameter r_0 in (4.10), since it controls the surface gravity κ of the black hole. Crucially, we cannot simply use the function $g(r)$ from the global Schwarzschild–AdS black hole metric, as the variation of r_0

will only allow us to explore a finite range of temperatures that may not include the deconfinement temperature. This is because the global Schwarzschild–AdS black hole branch of solutions has a minimum temperature which is not necessarily below the deconfinement temperature. We can overcome this difficulty by introducing another parameter corresponding to the value of T at the boundary, $T\big|_{y=1} = T_0$. Changing T_0 corresponds to changing the norm of the Euclidean time circle at the boundary, which in turn implies that the inverse temperature is modified as

$$\beta = \frac{2\,\pi\,\sqrt{T_0}}{\kappa}\,. \tag{4.15}$$

Note that we must also set $T = T_0$ on the AdS soliton end of the geometry, at $x = 1$.

To generate the data presented in the next section, we used 73852 grid points (equivalent to a resolution of 175×275 in the inner patch and 175×175 in the outer patch). The error in our numerical solution can be estimated by considering the maximum magnitude of the scalar $1 + R/\left[D\,(D+1)\right]$ on the grid, which suggests a relative error between 0.0001% and 0.01% for medium-sized plasma balls. Another useful measure of error is the maximum norm of the DeTurck vector $\sqrt{\xi^a \xi_a}$, which was found to have about the same order of magnitude as the Ricci scalar above. The highly deformed geometry of larger plasma balls naturally induces larger errors, and our numerical code typically fails when the error is estimated to be on the order of 1%.

4.3 Results

This section presents our main results. §4.3.1 characterizes the plasma ball geometry as we move along the family of solutions. In §4.3.2, we go on to study temporal Wilson loops, which are probes of confinement/deconfinement in the background of the plasma balls. Finally, in §4.3.3 we provide our results for the stress tensor of the dual field theory.

4.3.1 Characterizing the geometry of plasma balls

In this subsection, we characterize the horizon geometry of plasma balls. For small plasma balls (i.e. those at high temperatures), the dual black holes are much smaller than the radius of

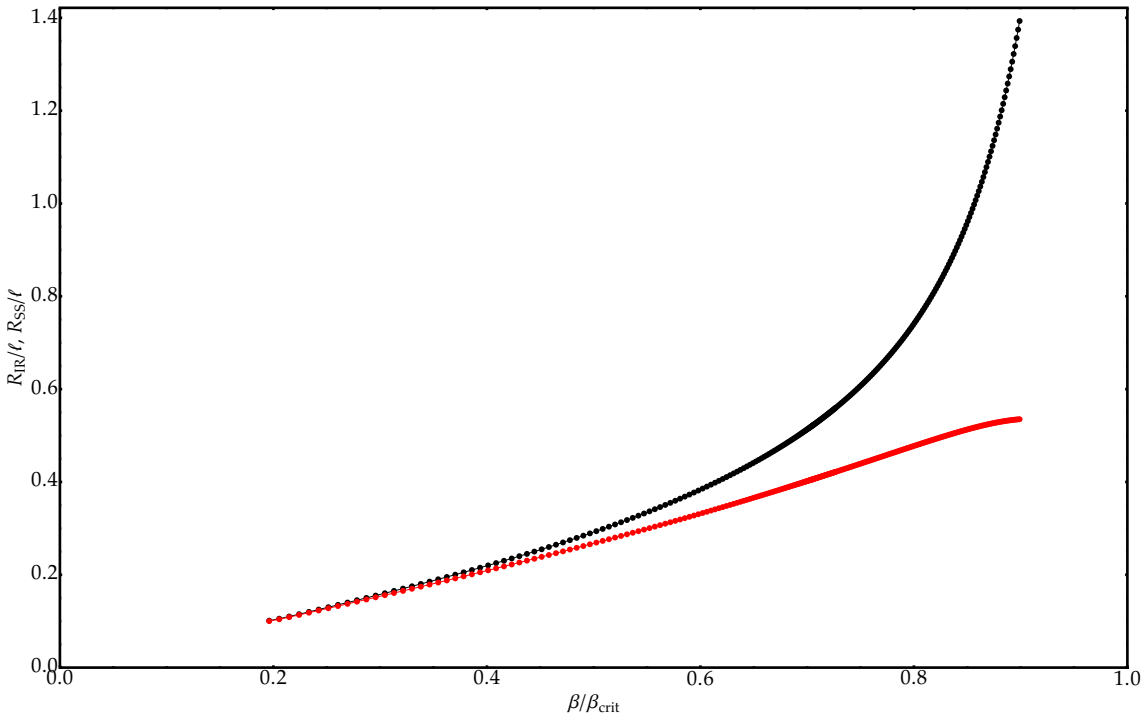

Figure 4.2: Plasma ball horizons can be characterized by two quantities: the radial extent $R_{\rm IR}$ (black) of the spatial S^2 along the IR bottom, and the radius $R_{\rm SS}$ (red) of the SS circle at the origin of the non-compact spatial dimensions. Here we plot these two quantities as functions of the scaled inverse temperature.

AdS and the size of the SS circle. As such, they look approximately like the asymptotically flat Schwarzschild solution, and the geometry of the spatial cross-sections of the horizon is roughly that of a round sphere. In the opposite limit, approaching the deconfinement temperature, plasma balls approaches the domain wall solutions of [42]. Therefore, we expect that the horizon will have finite extent in the SS direction at the IR bottom, while its size along the IR floor will diverge. Hence, for large plasma balls, the horizon of the dual black holes should look like a pancake extending along the IR bottom of the space-time.

To study the actual horizon shape, we can consider the induced metric on the spatial cross sections of the horizon:

$$\mathrm{d}s_H^2 = r_0^2 \left((1-a^2)^2\, R'\, \mathrm{d}\Omega_{(D-3)}^2 + a^2(2-a^2)\, S'\, \mathrm{d}\theta^2 + \frac{4\, B'}{2-a^2}\, \mathrm{d}a^2 \right) , \tag{4.16}$$

where the various functions are evaluated at $r = 0$ (the horizon). Recall that the topology of the horizon is S^{D-1}; we can measure the deformation of the horizon sphere by comparing the size of

the SS circle at the $a = 1$ equator and the size of the round S^{D-3} at the other equator, $a = 0$:

$$R_{\rm SS} = r_0 \sqrt{S'}\big|_{r=0,a=1}\,, \qquad R_{\rm IR} = r_0 \sqrt{R'}\big|_{r=0,a=0}\,. \tag{4.17}$$

For a round sphere, these two radii should be approximately equal, whereas for a pancaked black hole we should have $R_{\rm IR} \gg R_{\rm SS}$. Figure 4.2 displays these radii in the 5D case, measured in units of the AdS radius and plotted as a function of the inverse temperature. As this figure shows, in the high temperature limit both radii are equal to a very good approximation, and the black hole in the bulk is small compared to the radius of AdS. This demonstrates that the geometry of the horizon is indeed that of an approximately round sphere. As we lower the temperature toward the deconfinement temperature, we see that $R_{\rm SS}$ tends to saturate at a value of around 0.535. This suggests that in the limit $T \to T_d$, the extent of the horizon orthogonal to the IR bottom is finite. On the other hand, the size of the horizon along the IR bottom of the geometry is measured by $R_{\rm IR}$, which diverges at the deconfinement temperature. Of course, our expectation is that the bulk geometry of the plasma ball should approach the black brane solution as $T \to T_d$, where the SS circle has constant size everywhere. Given our choice of parameters, this means that $R_{\rm SS}$ should approach $1/2$ as $T \to T_d$. While one might suspect that our 'overshot' value of 0.535 is indicative of numerical error, our consideration of the dual stress tensor leads us to believe that this is not the case. Instead, this apparent saturation in the value of $R_{\rm SS}$ is actually a turning point beyond which it will decrease to $1/2$. More precisely, for the black brane family we know analytically that the (subtracted) holographic energy density decreases monotonically as $\beta \nearrow \beta_{\rm crit}$. However, in our solutions we find that the energy density at the centre of the plasma ball increases with β throughout the entire range for which our solution converged, in a similar fashion to R_{SS} (see Figure 4.7 in Section 4.3.3). For very large plasma balls the dependence on β of the energy density at the centre of the ball must become effectively brane-like, which implies the existence of a turning point in this dependence as $\beta \nearrow \beta_{\rm crit}$. It seems reasonable to expect that a corresponding turning point will exist in the $R_{\rm SS}$ dependence as well.

In Figure 4.3 we plot the horizon area, $\mathcal{A}_H$, as a function of radial extent of the horizon along the IR bottom, $R_{\rm IR}$, for 5D plasma balls. Notice the logarithmic scale of the x-axis in this plot. For small values of $R_{\rm IR}$, the area of the horizon scales like $R_{\rm IR}^3$, which is the expected behaviour for a spherically symmetric, static 5D Schwarzschild black hole. On the other hand, for large plasma

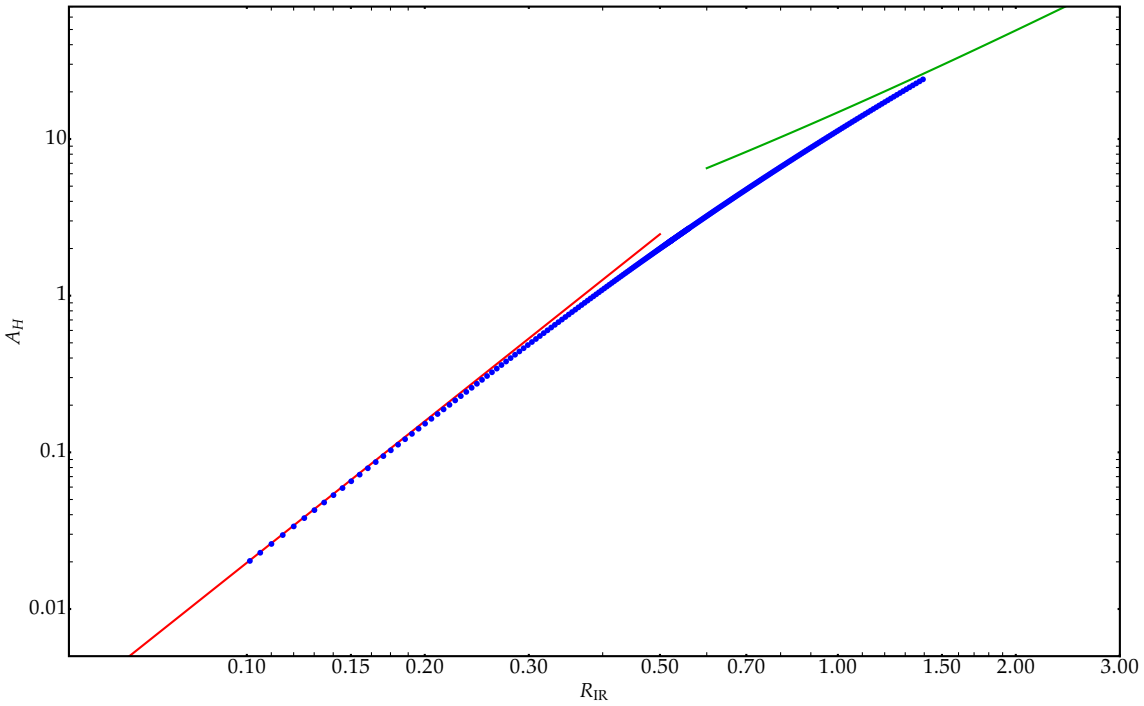

Figure 4.3: Horizon area A_H as a function of $R_{\rm IR}$, the radial extent of the horizon along the IR bottom, for 5D plasma balls. The red line shows the function $A_H = 2\pi^2 R_{\rm IR}^3$, which is the scaling behaviour of a perfectly-S^3 horizon. The green line shows the function $A_H = 2\pi^2 \left(R_{\rm IR}^2 R_{\rm SS} + R_{\rm IR} R_{\rm SS}^2\right)$, which is the scaling behaviour of a horizon that is only stretching in the spatial S^2 direction, while the SS circle maintains a constant size $R_{\rm SS} \equiv 1/2$. As $R_{\rm IR} \to \infty$ the behaviour becomes effectively $A_H \sim R_{\rm IR}^2$.

balls, the horizon area scales like $R_{\rm IR}^2$. This is the expected behaviour for a 5D black brane whose horizon is infinitely extended in two directions while the third one is compact. Our results therefore indicate that even though the plasma balls we have constructed are not parametrically much larger than the radius of AdS, they already exhibit some of the expected behaviours of a black brane. In fact, the Smarr relation in 5D implies that the mass of the large plasma balls scales with radius like $R_{\rm IR}^2$, just like the horizon area. Hence, our results confirm the predictions of [42]. We will provide additional evidence from the analysis of the stress tensor of the dual CFT.

Finally, to obtain a better intuition about the actual geometry of the horizon, we embed the horizon geometry (4.16), into Euclidean $\mathbb{R}^{D-1} \times S^1_{\rm SS}$ space,

$$\mathrm{d}s_E^2 = \mathrm{d}X^2 + \mathrm{d}Y^2 + X^2\,\mathrm{d}\Omega^2_{(D-3)} + Y^2\,\mathrm{d}\theta^2\,. \tag{4.18}$$

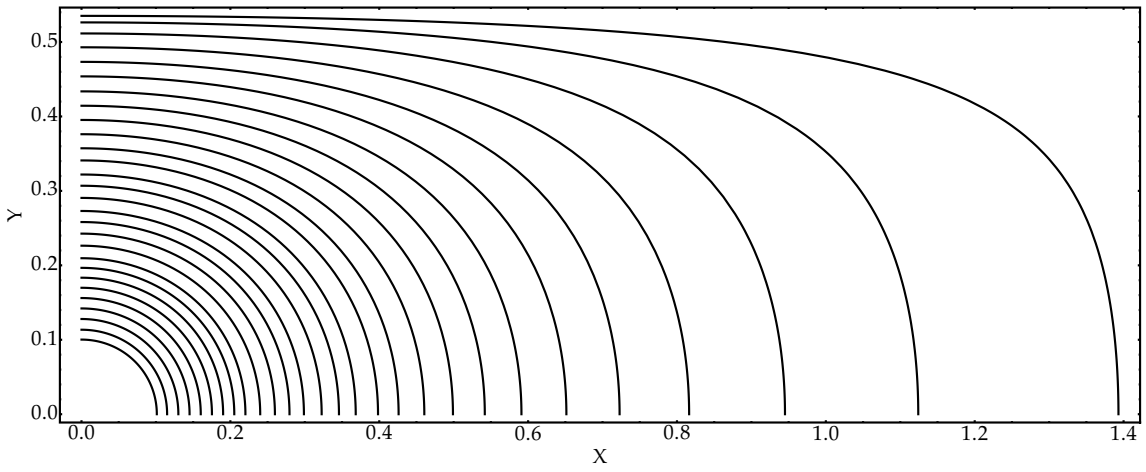

Figure 4.4: Embeddings of the horizon geometry into Euclidean $\mathbb{R}^{d-1} \times S^1_{\rm SS}$, for 5D solutions at $\beta/\beta_{\rm crit}$, roughly equally spaced between 0.2 and 0.9.

Figure 4.4 shows plots of Y vs. X for some of the solutions that we have constructed. As expected, this demonstrates that small plasma balls appear to be round, while large plasma balls have a pancake-like shape.

4.3.2 Wilson loops as a probe of confinement

The standard order parameter for deconfinement is the expectation value of a Wilson loop wrapping the Euclidean time circle:

$$\langle |\mathrm{tr}\, W| \rangle = \left\langle \left| \frac{1}{N} \mathrm{tr} \left(P \exp \left[i \oint_{\mathcal{C}} A_\tau \, \mathrm{d}\tau \right] \right) \right| \right\rangle , \tag{4.19}$$

where A_μ the gauge field, and 'P exp' denotes the path-ordered exponential. The temporal Wilson loop measures the cost in free energy of perturbing the system by an external quark. If the system is in a confining state, then the cost in free energy is infinite and $\langle |\mathrm{tr}\, W| \rangle = 0$, whereas in a deconfined state, the cost in free energy is finite and the temporal Wilson loop has a finite expectation value [41].

Refs. [95, 96] provide a prescription for calculating Wilson loops in AdS/CFT. According to this prescription, one should consider the action of a classical string wrapping the Euclidean time circle on the boundary in a contour $\mathcal{C} = P \times S^1_\beta$, where P is a point in the transverse directions, and

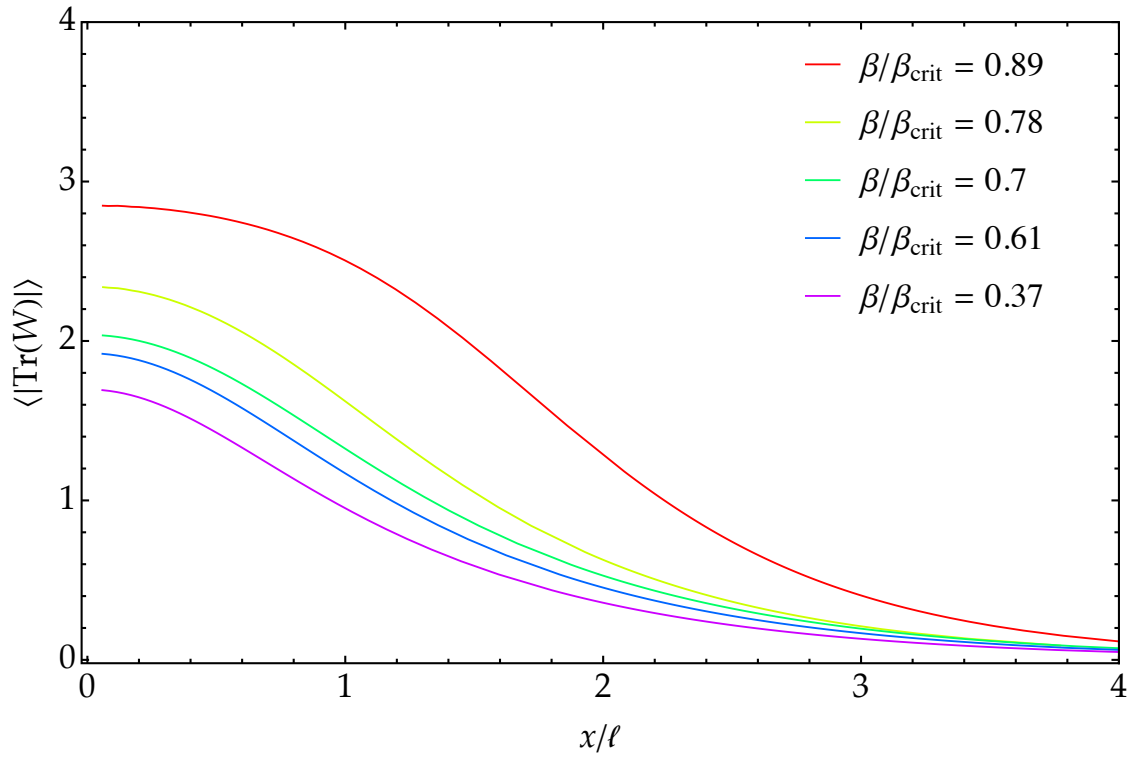

Figure 4.5: Expectation value of temporal Wilson loops, $\langle|\text{tr}\, W|\rangle$, for different temperatures in the 5D bulk case. For temperatures closer to the deconfinement temperature, the black holes in the bulk are larger and hence $\langle|\text{tr}\, W|\rangle$ is non-zero over a larger region of the boundary. The size of this region provides an estimate of the actual size of the plasma ball.

smoothly extending into the bulk,

$$\langle|\text{tr}\, W|\rangle \sim e^{-S_{\text{string}}}\,. \tag{4.20}$$

This classical string action (i.e. the area of a minimal surface) diverges, and its divergence is proportional to the length of the circumference $\mathcal{C}$ on the boundary. By subtracting the universal infinite piece, we can define a regularised classical string action, which can be meaningfully compared with the CFT predictions. It is this regularised area that we calculate below.

When the Euclidean time circle is non-contractible in the bulk, no classical string world-sheet ending on $\mathcal{C}$ exists, and hence $\langle|\text{tr}\, W|\rangle \simeq 0$ [41]. This is the expected result for the confining state, whose bulk dual is the AdS soliton. On the other hand, if the Euclidean time circle is contractible, then a minimal surface ending on $\mathcal{C}$ which smoothly caps off in the bulk does exist, and hence $\langle|\text{tr}\, W|\rangle \neq 0$. This is the deconfined phase and its dual geometry is the black brane. For our

plasma ball geometries, the Euclidean time circle is contractible in the bulk, because at any finite temperature there is a finite size horizon sitting at the IR bottom of the geometry. Therefore, we expect that a minimal surface ending on $\mathcal{C}$ should exist and so the temporal Wilson loop should have a non-zero expectation value. Indeed, from the boundary CFT perspective, there is a finite region of deconfined plasma and the temporal Wilson loop should be sensitive to it. Since $\langle|\mathrm{tr}\, W|\rangle$ is a probe of deconfinement, the expectation value of this operator should provide a measure of the size of the deconfined region. In other words, if the size of the Wilson loop on the boundary is much larger than the size of the deconfined region, then the corresponding expectation value should be exponentially small. By contrast, if the size of the Wilson loop is of the same order as the size of the deconfined region, then the Wilson loop should have a finite expectation value.

To see this, we consider temporal Wilson loops ending at some fixed value of the radial coordinate x along the boundary (4.9). We therefore have $\mathcal{C} = x \times S^1_\beta$ for our Wilson loops. In Figure 4.5, we display $\langle|\mathrm{tr}\, W|\rangle$ as a function of the radial co-ordinate on the boundary, x, for plasma balls of different sizes and therefore different temperatures. As the figure shows, when the size x of the Wilson loop on the boundary is comparable to the size of the bulk horizon along the IR bottom, the expectation value of the Wilson loop is non-zero. In fact, the typical size of a plasma ball can be estimated by the value of x for which $\langle|\mathrm{tr}\, W|\rangle$ has decayed by one e-fold with respect to the corresponding value at the centre of the ball. The estimate we obtain using this simple method is comparable to the size of the black hole along the IR bottom, R_{IR}. Therefore, we conclude that the expectation values of temporal Wilson loops offer a measure of the size of the deconfined region on the boundary.

4.3.3 Dual stress tensor

In this final subsection, we study the vacuum expectation value of the stress tensor of the plasma balls. Having constructed plasma ball solutions numerically, we extract the stress tensor using standard holographic renormalization techniques [40].

The prescription of [40] normalises the stress tensor such that it vanishes when the bulk geometry is AdS. The AdS soliton would therefore have a non-zero vacuum expectation value, which can be interpreted as a Casimir energy. Since the AdS soliton is the lowest energy state obeying our

boundary conditions [91], in the following analysis we subtract the stress tensor of the soliton, so that the final stress tensor vanishes on the AdS soliton background. We obtain,

$$\langle T^{\rm sub}_{ij}\rangle\, \mathrm{d}x^i\, \mathrm{d}x^j = c_{\rm eff}\, D \Bigg[-\frac{t^{(D)}(x)}{T_0}\, \mathrm{d}\tau^2 + s^{(D)}(x)\, \mathrm{d}\theta^2 + \frac{k_x^2\, x^2\, r^{(D)}(x)}{(1-x^2)^2}\, \mathrm{d}\Omega^2_{(D-3)} + \frac{k_x^2 (1+x^2)^2}{(1-x^2)^4} \left(\frac{t^{(D)}(x)}{T_0} - s^{(D)}(x) - (D-3)\, r^{(D)}(x) \right) \mathrm{d}x^2 \Bigg]. \tag{4.21}$$

where $c_{\rm eff}$ is the effective central charge [90],

$$c_{\rm eff} = \frac{\ell^{D+1}}{16\,\pi\, G_N^{D+1}}\,. \tag{4.22}$$

Note that (4.21) is manifestly traceless. As usual, this follows from solving the bulk equations of motion near the boundary. In (4.21), the functions $t^{(D)}(x)$, $s^{(D)}(x)$ and $r^{(D)}(x)$ denote the coefficients of the order D terms in the near boundary expansions of $T(x,y)$, $S(x,y)$ and $R(x,y)$, in the far-region metric (4.9) and in our working coordinates:

$$\begin{aligned} T(x,y) &= T_0 + (1-y^2)^D\, t^{(D)}(x) + \ldots \\ S(x,y) &= 1 + (1-y^2)^D\, s^{(D)}(x) + \ldots \\ R(x,y) &= 1 + (1-y^2)^D\, r^{(D)}(x) + \ldots \end{aligned} \tag{4.23}$$

We can extract these functions from our numerical solutions by fitting our data to this near boundary behaviour. Of course, the functions are not independent; conservation of the stress energy tensor imposes one differential constraint among them. As is well-known, the conservation of the dual stress energy tensor is guaranteed by the bulk equations of motion. Therefore, we need not impose this constraint, but we can use it to asses the numerical error in the stress energy tensor. In the data below, we estimate the error to be less than 5%.

Figure 4.6 shows the various components of the stress tensor for a five-dimensional plasma ball with temperature $\beta/\beta_{\rm crit} = 0.90$, as a function of the radial coordinate along the boundary x. This particular plasma ball is quite large, with $R_{\rm IR}/\ell = 1.35$. Therefore, it is unsurprising that the stress tensor exhibits features which are qualitatively similar to the stress tensor of the domain wall solution ([42], see also [90]). In particular, at the centre of the ball, $\langle T^{\rm sub}_{\tau\tau}\rangle$ and $\langle T^{\rm sub}_{\theta\theta}\rangle$ have a similar magnitude and opposite sign. Note also that the components of the stress tensor along

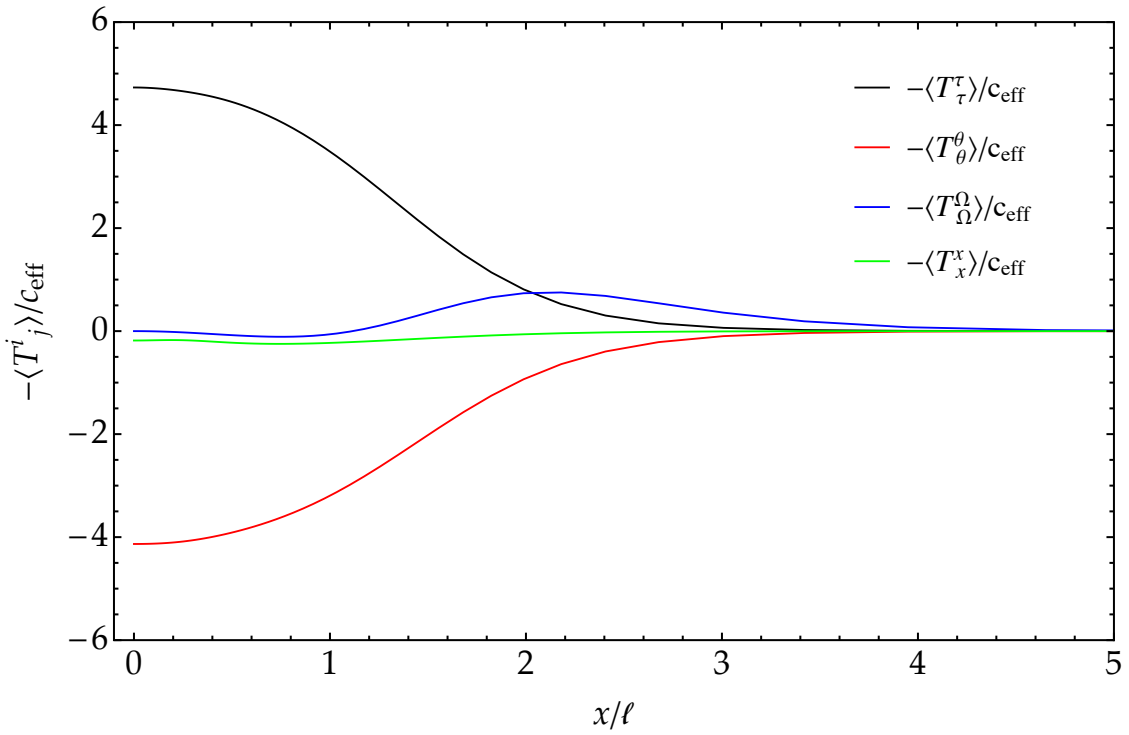

Figure 4.6: Components of the subtracted stress tensor for a 5D plasma ball with temperature $\beta/\beta_{\rm crit} = 0.90$. This stress tensor shows some of the qualitative features of the stress tensor of the domain wall solution [42]. In particular we find that the tension of the wall, measured by tangent components of the stress tensor $-\langle T^\Omega_\Omega\rangle$, is positive.

the sphere directions (the tangential components) obey $\langle T^{\rm sub}_{\Omega\Omega}\rangle < 0$ throughout the boundary of the ball, resulting in a positive tension.

Figure 4.7 displays the subtracted stress tensor components $\langle T^r_{\ r}\rangle$ (red) and $-\langle T^\tau_{\ \tau}\rangle$ (blue). These correspond to the radial pressure and the energy density at the centre of the plasma ball respectively. The pale solid lines in Figure 4.7 show the same components of the stress tensor dual to a black brane, at the same range of temperatures. As $\beta \nearrow \beta_{\rm crit}$, the plasma ball curve should essentially behave like the black brane curve, and indeed we clearly observe a turning point in the radial pressure toward the end of our available dataset. In order to see a turning point in the energy density, it appears we would need to construct larger plasma balls. However, the fact that the dotted blue curve must asymptote to the solid pale blue curve as $\beta \nearrow \beta_{\rm crit}$ suggests that a turning point in the energy density should exist. Our results also imply that for large, finite plasma balls, the energy density at the centre of the ball is greater than the energy density for the black brane at the deconfinement temperature; supporting the conjecture in [42].

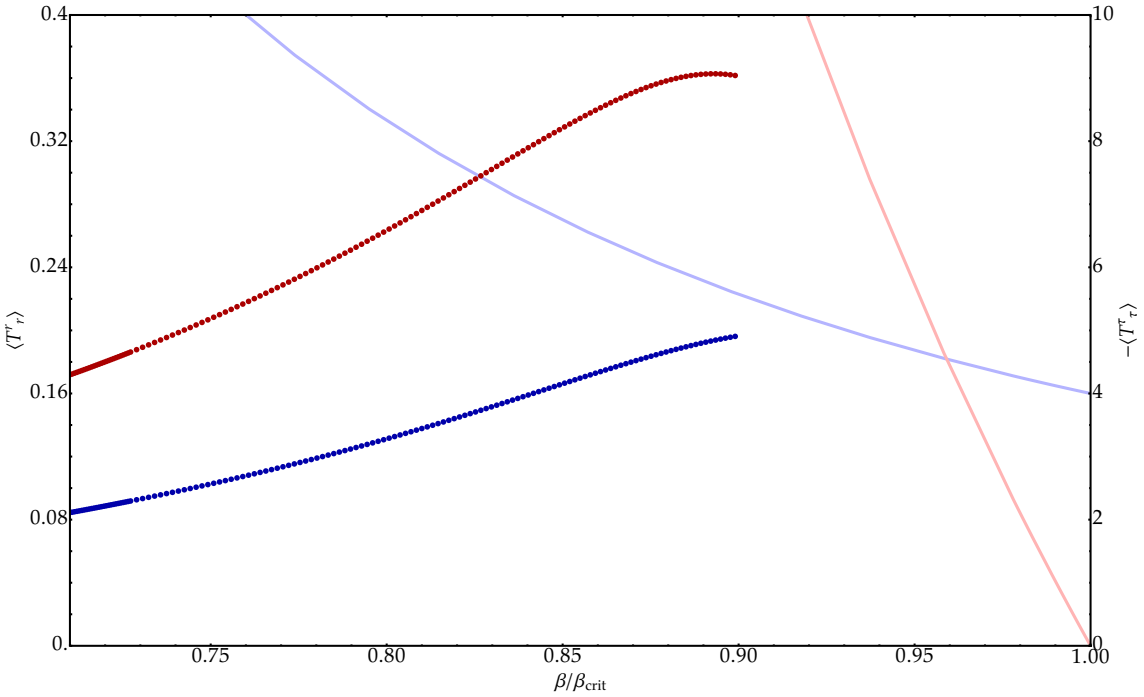

Figure 4.7: The subtracted stress tensor components $\langle T^r{}_r \rangle$ (red) and $-\langle T^\tau{}_\tau \rangle$ (blue) at the centre of the plasma ball, as a function of the inverse temperature $\beta/\beta_{\rm crit}$. Pale solid lines show the corresponding components of the stress tensor dual to the black brane.

We can assess how close our large plasma balls are to the domain wall solution by investigating the tension of the ball in more detail. [42] computed the tension in the domain wall limit and found,

$$\mu_{D-1} = C_D, \epsilon_c \left(\frac{L}{2\pi}\right)^2 \qquad C_4 = 2.0\,, \qquad C_5 = 1.7\,, \tag{4.24}$$

where the constants $C_{4,5}$ are determined numerically. For very large plasma balls of size $\rho \gg L$, the interior of the ball should be well-described by the homogeneous deconfined plasma phase, and in the vicinity of the edge these solutions should reduce to a domain wall. One can infer the behaviour of these large plasma balls by balancing the pressure of the internal deconfined phase with the tension of the wall, which yields

$$P = (D-3)\frac{\mu}{\rho}\,. \tag{4.25}$$

Since we know the pressure P of the deconfined phase as a function of the temperature, we can use (4.25) to obtain the corrections to temperature above the deconfinement transition, as a function of

the size ρ of the ball [90],

$$T = \frac{1}{L}\left(1 + \frac{(D-3)\,C_D}{(2\pi)^2}\frac{L}{\rho} + O\left(\frac{L^2}{\rho^2}\right)\right). \tag{4.26}$$

Using this relation, we can then independently estimate the values of C_4 and C_5, for comparison with the values reported in [42], by considering data from our largest plasma balls. Our best estimates for these parameters are

$$C_4 = 1.9\,, \qquad C_5 = 1.8\,. \tag{4.27}$$

4.4 Summary

In this chapter, we numerically constructed localized spherical black hole solutions in 5D and 6D in the confining background of the AdS soliton. From the dual CFT perspective, our solutions correspond to localized balls of deconfined plasma surrounded by the confining vacuum. These solutions are parametrized by temperature, and exist above the deconfinement temperature.

At temperatures that are much higher than the deconfinement temperature, these black holes are small relative to the radius of AdS, and are approximately spherically symmetric. By contrast, as the deconfinement temperature is approached, the extent of the horizon along the IR bottom of the geometry grows without bound, and the horizon adopts a pancake-like shape. In the region near the edges, these large black holes look like domain walls, similar to the one constructed in [42] but at a different temperature. In the region far from the edges, the geometry of the space-time is well-approximated by that of a homogeneous black brane at the corresponding temperature. As [42] predicted, and as we have confirmed, these localized black holes have the unusual property that in the infinite mass limit they have constant temperature (namely the deconfinement temperature).

The black holes we have constructed are thermodynamically unstable, since they have negative specific heat (they become cooler as their mass increases). At the deconfinement temperature their mass should become infinite. We have computed the spectrum of the Lichnerowicz operator around our solutions, and we find that it has a single negative mode, with no zero modes appearing

along the branch of solutions. This suggests that these black holes are classically stable under perturbations that preserve the spatial symmetries of these backgrounds. However, in the full quantum theory these black holes should be unstable to evaporation through the emission of Hawking radiation. The absolute value of the negative mode becomes smaller as we approach the deconfinement temperature; this makes it computationally challenging to construct very large black holes, as the operator we must invert at each Newton step has a near-zero mode. Although we do not address it here, it would be interesting to perform a detailed analysis of the quasi-normal mode spectrum of these black holes, as it could provide insight into near-equilibrium plasma physics in confining backgrounds. We have also computed temporal Wilson loops in the background of the plasma balls, and have verified that their expectation value provides an estimate of the size of the deconfined region.

CHAPTER 5

Black Rings in Global AdS

The work presented in this chapter was published as [97]. The paper was co-authored with Pau Figueras. The numerical setup, computer code, and analysis are entirely my own work. The majority of the numerical solutions produced are also my own work, however a small number of data points were generated by Figueras.

5.1 Introduction

The study of equilibrium black holes in $D \geq 5$ has revealed that the physics of these objects can be very different from that of their four-dimensional counterparts. In particular, there exist black holes with a non-spherical horizon topology, such as black rings [14, 15, 98], black ringoids [99], and regular multi-black hole spacetimes in vacuum [16, 17, 18, 19, 20], among others [100, 101, 102, 103]. These types of black objects do not exist in $D = 4$, and thus they possess fundamentally new physical properties. A review on this topic can be found in [104]. A natural question to ask is whether these objects can also exist in asymptotically anti-de Sitter (AdS) space. This question is further motivated by the gauge/gravity correspondence [24], according to which new types of stationary black holes should correspond to new finite temperature phases of gauge theories. Furthermore, AdS is a maximally symmetric space and, as such, can be regarded from a mathematical perspective as being as fundamental as Minkowski space. The study of black holes in AdS is therefore also an interesting question in its own right.

Through the AdS/CFT duality, one can look for new types of black holes by solving the equations

of motion of the dual field in their *hydrodynamic regime*, in which they simplify considerably. Using this approach, [105] looked for solutions to the Navier–Stokes equations on the Einstein static universe, $\mathbb{R}_t \times S^3$, corresponding to rotating fluid configurations. However, they found only stationary fluid configurations which are dual to known rotating spherical black hole solutions in AdS [106, 43, 107]. Still, the assumption of AdS asymptotics also allows for more general boundary conditions, such as Scherk–Schwarz compactifications of AdS. In these settings, [108] constructed solutions to the relativistic Navier–Stokes equations corresponding to rotating plasma balls and plasma rings; hence the authors were able to study the phase diagram of rotating black holes in such spacetimes. Of course, this approach can only capture the physics of black holes that admit a hydrodynamic limit, which is unfortunately not always the case.

This chapter concerns the numerical construction of stationary black rings in global AdS_5. This choice of dimensionality is particularly interesting from the point of view of the gauge/gravity correspondence. These black rings will have horizons with topology $S^1 \times S^2$. Techniques for generating analytic solutions [26, 109, 110] that have been successful in a 5D asymptotically flat (AF) setting do not seem to straightforwardly extend to AdS. Nevertheless, an approximate solution for black rings in global AdS was given in [111].

In the AF setting, black rings are classified as either *thin* or *fat*, depending on the ratio between the size of the horizon's S^1 (denoted R_{S^1}) and S^2 (denoted R_{S^2}). By working in AdS, the AdS radius ℓ provides yet another lengthscale which can influence the physics of black rings. In order to take all three lengthscales into account, the following terminology was introduced in [111] to classify black rings in AdS:

- *Thin* rings have $R_{S^2} \ll R_{S^1}$, while *fat* rings have $R_{S^2} \gtrsim R_{S^1}$.
- *Small* rings have $R_{S^2} < \ell$, while *large* rings have $R_{S^2} > \ell$.
- *Short* rings have $R_{S^1} < \ell$, while *long* rings have $R_{S^1} > \ell$.

According to this classification, the approximate method of [111] is valid for rings which are small and thin, regardless of whether they are short or long. Based on quantitative results for thin rings and some educated guesses, [111] proposed a phase diagram of AdS black rings in

the microcanonical ensemble which is qualitatively similar to the AF one, but with the angular momentum compressed into the range $J \leq M\,\ell$.

By using numerical methods, we can access all AdS black rings: thin or fat, small or large, short or long. This allows us to provide a complete phase diagram of 5D black rings in AdS. In particular, we confirm that fat rings merge with the spherical black holes at a singular solution with zero area, as in the AF case. This suggests that the pattern of connections between various stationary black hole phases of different topologies conjectured in [111] is indeed correct in $D \geq 6$. One of the main unanswered questions in [111] was whether rings that are both thin and large exist. We address this question here and, quite confidently, find no evidence for such thin large rings.

The rest of this chapter is organised as follows. In §5.2 we describe our numerical construction of black rings in AdS and in §5.3 we explain how we calculate various physical quantities, in particular the mass, from our numerical solutions. In §5.4, we study the geometry of the horizon of the AdS black rings, and in §5.5 we study the thermodynamics and produce phase diagrams for the grand canonical and microcanonical ensembles. Finally, in §5.6, we extract the v.e.v. of the stress tensors of the dual CFT and compare them with that of the spherical black holes.

Throughout this chapter, we use a tilde to denote quantities which have been nondimensionalised with respect to the AdS radius ℓ. In particular, we define the following quantities

$$\tilde{M} := M\,\ell^{-2} \quad , \quad \tilde{\kappa} := \kappa\,\ell \quad , \quad \tilde{A}_H := A_H\,\ell^{-3} \quad , \quad \tilde{\Omega}_H := \Omega_H\,\ell \quad , \quad \tilde{J} := J\,\ell^{-3}. \tag{5.1}$$

5.2 Numerical construction of AdS black rings

This section provides details of the method used to numerically construct black rings in global AdS. We begin in §5.2.1 by discussing our choice of coordinates and the formulation of our spacetime metric ansatz. We then proceed to explain our construction of the reference metric in §5.2.2, followed by a description of our boundary conditions in §5.2.3. We conclude the section with further technical details in §5.2.4

5.2.1 Basic setup

Our goal is to construct black rings in 5D with the asymptotics of global AdS. These are stationary black hole solutions to the vacuum Einstein equations with a negative cosmological constant, whose spatial horizon topology is $S^1 \times S^2$. For simplicity, we shall restrict ourselves to *singly-spinning* black rings, i.e. those which are only rotating along the S^1 direction. Our solutions are therefore the AdS generalisation of the black rings of [14].

The isometry group of the AdS rings is the same as that of the AF rings, namely $\mathbb{R}_t \times U(1)^2$, corresponding to time translations and rotations on the two 2-planes. In the AF case, the Einstein equations restricted to spacetimes with this isometry group are completely integrable. This makes it possible to explicitly construct all solutions within this class using an algebraic procedure [26, 109, 110]. Unfortunately, it is not currently known whether this integrability persists in the presence of a cosmological constant. Therefore, we shall rely here on numerical methods to construct our solutions.

The black ring metric can be written in a C-metric type of coordinates [14, 112, 98], which are well-adapted to its horizon geometry. Specifically, they foliate spatial slices of the spacetime with surfaces of ring-like topology. Similar coordinates can be used in AdS [113]. The main drawback of this approach is that spatial infinity is represented by a single point in this coordinate system. In order to use this in numerical work, it would be necessary to first analyse the singular behaviour of the various metric components near the boundary of AdS. Furthermore, spatial infinity would be very poorly resolved when discretised, and this could become an issue when it comes to extracting the stress tensor of the dual CFT. For this work, we instead cover the spacetime with two coordinate patches: one is adapted to the *"outer"* region near the boundary of AdS, whilst the other is adapted to the *"inner"* region near the horizon of the ring.

The geometry near the horizon of the AdS black ring can be treated as a deformation of the near-horizon geometry of the AF black ring. This deformation will not be small in general, but at very high temperatures, where the rings are much smaller than the AdS radius, we can expect the near-horizon geometry to be very close to that of AF rings. Therefore, we build our inner region

ansatz by “dressing" the AF black ring metric:

$$\begin{aligned} \mathrm{d}s^2_{\text{inner}} =& -\mathcal{T}_0\, e^{\mathcal{T}}\, \mathrm{d}t^2 + \mathcal{X}_0\, e^{\mathcal{X}}\, \mathrm{d}x^2 + \mathcal{Y}_0\, e^{\mathcal{Y}} (\mathrm{d}y - \mathcal{W}\, \mathrm{d}x)^2 \\ &+ \mathcal{U}_0\, e^{\mathcal{U}}\, \mathrm{d}\phi^2 + \mathcal{V}_0\, e^{\mathcal{V}}\, (\mathrm{d}\psi - \mathcal{Z}_0\, (1+\mathcal{Z})\, \mathrm{d}t)^2 \ , \end{aligned} \tag{5.2}$$

where $\mathcal{F} \coloneqq (\mathcal{T}, \mathcal{X}, \mathcal{Y}, \mathcal{U}, \mathcal{V}, \mathcal{W}, \mathcal{Z})$ are unknown functions of (x, y), and $\mathcal{F}_0 \coloneqq (\mathcal{T}_0, \mathcal{X}_0, \mathcal{Y}_0, \mathcal{U}_0, \mathcal{V}_0, \mathcal{Z}_0)$ are functions which are analytically prescribed so that we recover the AF black ring line element when $\mathcal{F} \equiv 0$. In order to normalise the coordinate ranges and impose boundary conditions, we transform the ring-like (x, y) coordinates from the ones described in [112, 98] to

$$x \to \cos(\pi\, x)\,, \qquad y \to -\frac{1+\nu+(1-y)\cos(\pi\, y)}{2\,\nu}\,. \tag{5.3}$$

The ranges of these transformed coordinates then become $0 \leq \{x, y\} \leq 1$ in this near horizon region. Here ν is the dimensionless parameter introduced in [112], which is related to the ring's surface gravity κ and horizon angular velocity Ω_H by

$$\nu = \frac{\Omega_H}{\sqrt{4\,\kappa^2 + \Omega_H^2}}. \tag{5.4}$$

For the AF black ring the range of this parameter is $0 < \nu < 1$. However, in AdS we found that the lower bound increases as the temperature lowers (see §5.5.1). Writing down the AF ring in these new (x, y) coordinates allows us to identify the expressions for the functions $\mathcal{F}_0$.

Near the conformal boundary, we expect the black ring spacetime to be a small deformation of pure global AdS. Therefore, the metric in the outer region is written in a form which is manifestly asymptotically AdS. The most general one which is closed under diffeomorphisms preserving the $\mathbb{R}_t \times U(1)^2$ isometry is given by

$$\begin{aligned} \mathrm{d}s^2_{\text{outer}} =& -\left(1+\lambda^2\, R^2\right)\, e^T\, \mathrm{d}t^2 + \frac{e^X}{1+\lambda^2\, R^2}\, (\mathrm{d}R - W\, \mathrm{d}a)^2 \\ &+ R^2 \left[\tfrac{\pi^2}{4}\, e^Y\, \mathrm{d}a^2 + \cos^2\!\left(\tfrac{\pi\, a}{2}\right) e^U\, \mathrm{d}\phi^2 + \sin^2\!\left(\tfrac{\pi\, a}{2}\right) e^V\, (\mathrm{d}\psi - \mathcal{Z}_0\, (1+Z)\, \mathrm{d}t)^2 \right], \end{aligned} \tag{5.5}$$

where $\lambda \coloneqq \ell^{-1}$ is the inverse of the AdS radius, $F \coloneqq (T, X, Y, U, V, W, Z)$ are unknown functions of (R, a), and $\mathcal{Z}_0$ is identical to the function that appeared in (5.2) but transformed into (R, a) coordinates via a relation which we explain in due course. We use λ instead of ℓ as a parameter in

our ansatz to allow us to easily connect to AF rings by setting $\lambda = 0$, rather than trying to make ℓ very large.

In order to cover the AdS boundary, we define a compact radial coordinate r via

$$R = \frac{r}{1 - r^2/k^2}\,, \tag{5.6}$$

where k is some constant lengthscale. The ranges of these coordinates are $r_{\min} \leq r \leq k$ and $0 \leq a \leq 1$, where $r_{\min}$ is a parameter that sets the location where we switch between the outer and inner region patches, which is an arbitrary gauge choice within the limits described below. For this work, we take $r_{\min} \approx k/2$.

The outer region coordinates (r, a) are related to the inner region ones (x, y) by a simple coordinate transformation,

$$x = (1 - r/k)\,\cos\left(\tfrac{\pi a}{2}\right)\,, \qquad y = 1 - (1 - r/k)\,\sin\left(\tfrac{\pi a}{2}\right)\,. \tag{5.7}$$

The transformation becomes singular when $r = 0$, which sets a lower bound on the parameter $r_{\min}$. We choose to fix $k = 1$ throughout, thereby setting the scale for the solutions. A graphical depiction of the two coordinate patches is given in Figure 5.1.

5.2.2 Reference metric

We follow [45] and construct a suitable reference metric $\bar{g}$ on the manifold $\mathcal{M}$ by considering

$$\bar{g}_{\mu\nu}\,\mathrm{d}x^\mu\,\mathrm{d}x^\nu = [1 - I(\hat{r})]\;\bar{g}^{\text{outer}}_{\mu\nu}\,\mathrm{d}x^\mu\,\mathrm{d}x^\nu + I(\hat{r})\;\bar{g}^{\text{inner}}_{\mu\nu}\,\mathrm{d}x^\mu\,\mathrm{d}x^\nu\,, \tag{5.8}$$

where $\bar{g}^{\text{outer}}_{\mu\nu}$ is obtained from (5.5) and $\bar{g}^{\text{inner}}_{\mu\nu}$ is obtained from (5.2) after setting $\mathcal{F} \equiv 0 \equiv F$. $I(\hat{r})$ is an interpolating function which is a function of a suitably defined coordinate $\hat{r}$ satisfying $\hat{r} = 0$ at the horizon and $\hat{r} = 1$ at infinity. The function $I(\hat{r})$ must be smooth and must lie within the range $0 \leq I(\hat{r}) \leq 1$, with equality only at the endpoints. We found that the system behaves better if we further require $I(\hat{r})$ to be monotonic. Previous works, e.g. [45], used an interpolating function with a compact support which is limited to a subregion that did not extend to the boundaries.

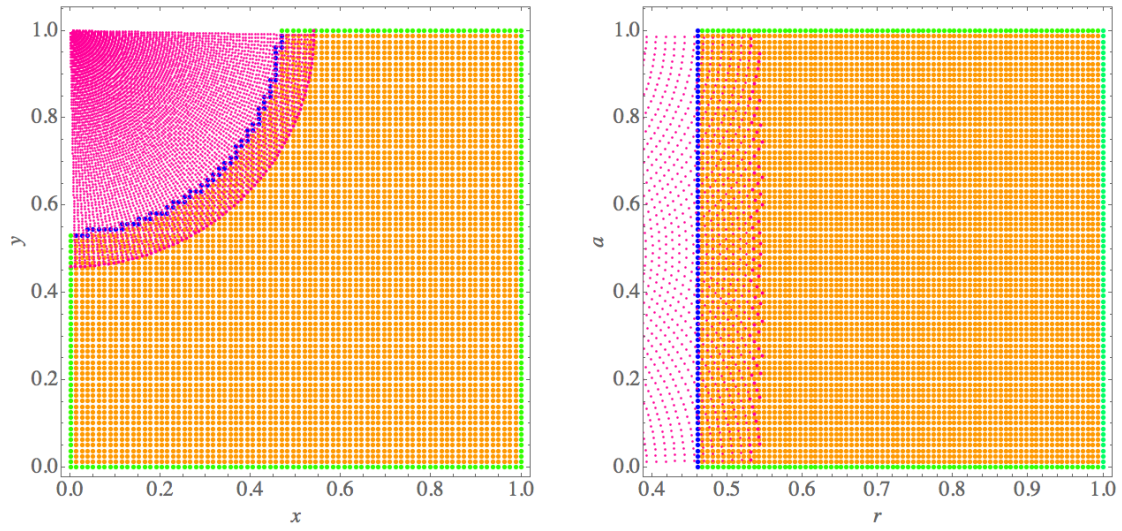

Figure 5.1: Our computational grid, shown in the inner (left) and outer (right) region coordinates. Orange dots show grid points where the inner patch ansatz (5.2) is used. Pink dots show grid points where the outer patch ansatz (5.5) is used. Green dots show grid points where boundary conditions are imposed, and blue dots show grid points where function values are obtained by interpolating data from the other patch. Note that these diagrams are only illustrative, as the grid resolution used for actual calculations is significantly higher than shown above.

However, such a function tends to have large derivatives which are inevitably inherited by the solutions. For the AdS black ring setup, we found that this proved problematic for the numerics. Instead, we use an interpolation function which *is* supported near the boundaries, but whose normal derivatives vanish up to at least fourth order, both at the horizon and at the boundary of AdS. This still ensures that the reference metric (5.8) satisfies all the boundary conditions without introducing excessively large gradients into various functions. To construct such an interpolating function, we exploit the fact that the (x, y) coordinates (5.3) can in fact cover our entire computational domain. This allows us to globally define a non-compact "radial" coordinate, centered at the AdS boundary $(x, y) = (0, 1)$, by

$$\hat{R} := \sqrt{\left(\frac{x}{1-x}\right)^2 + \left(\frac{1-y}{y}\right)^2}. \tag{5.9}$$

$\hat{R}$ can be compactified to obtain $\hat{r} \in [0, 1]$ via

$$\hat{r} := \frac{1}{1+\hat{R}}. \tag{5.10}$$

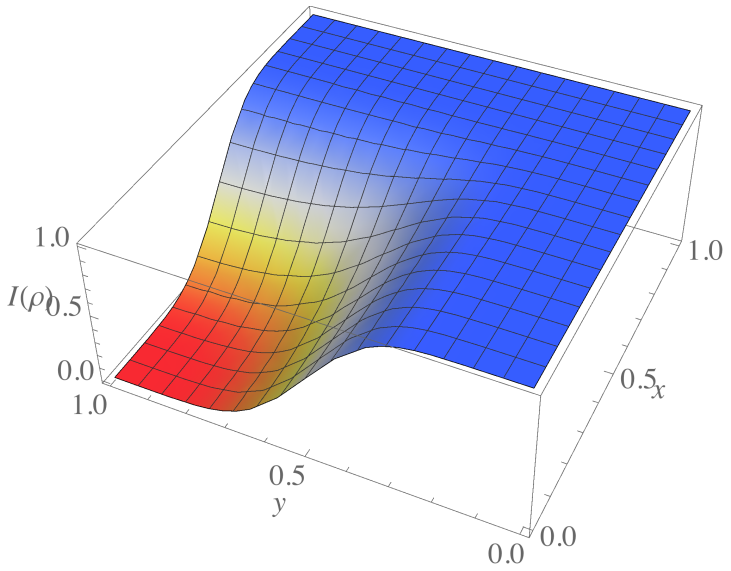

Figure 5.2: The interpolating function $I(\hat{r})$ in terms of the inner region coordinates (x, y). This is a smooth function defined everywhere, but it is not compactly supported.

Crucially, this gives us $y \sim \hat{r}$ as $\hat{r} \to 0$ in the inner patch, and $r \sim \hat{r}$ as $\hat{r} \to 1$ in the outer patch. In terms of $\hat{r}$, we may now define our interpolating function as

$$I(\hat{r}) := 1 - \hat{\varrho}^4 \left(6 - 8\,\hat{\varrho} + 3\,\hat{\varrho}^2\right)^2 \quad \text{where} \quad \hat{\varrho} := \sin^2\left(\frac{\pi\,\hat{r}}{2}\right). \tag{5.11}$$

Note that, by defining I in terms of $\hat{\varrho}$ rather than $\hat{r}$, we ensure that *all* even-order normal derivatives of I vanish at $y = 0$ and $r = 1$. We depict our choice of interpolating function in Figure 5.2.

5.2.3 Boundary conditions

The boundary conditions imposed on our unknown functions are detailed below. At all boundaries, the reference metric $\bar{g}$ must also satisfy the same conditions. Our construction above ensures that this is the case.

Near region patch:

- Horizon ($y = 0$): regularity of the spacetime metric requires that we impose a Neumann boundary condition $\partial_y \mathcal{F}\big|_{y=0} = 0$ on all functions, except for $\mathcal{W}$ which has to vanish. To

ensure that the temperature and angular velocity of the AdS rings are the same as those of the AF ring, we further impose $\mathcal{Y} = \mathcal{T}$ and $\mathcal{Z} = 0$ on this boundary.

- S^1 axis ($y = 1$): regularity of the spacetime metric requires that we impose a Neumann boundary condition $\partial_y \mathcal{F}\big|_{y=1} = 0$ on all functions, except for $\mathcal{W}$ which has to vanish. To avoid conical singularities, we further require $\mathcal{Y} = \mathcal{V}$ on this boundary.

- S^2 axes ($x = 0$ and $x = 1$): again, regularity of the spacetime metric requires that we impose a Neumann boundary condition $\partial_x \mathcal{F}\big|_{x=0,1} = 0$ on all functions, except for $\mathcal{W}$ which has to vanish. To avoid conical singularities, we further require $\mathcal{X} = \mathcal{U}$ on these boundaries.

- Interpolation boundary (blue dots in Fig. 5.1 (left)): the value of each function is determined by first interpolating the values of the functions in the outer patch, then applying the coordinate transformations.

Far region patch:

- Spacelike infinity ($r = 1$): our ansatz is manifestly asymptotically AdS, provided that we impose a Dirichlet boundary condition $F = 0$ on all functions.

- S^3 axes ($a = 0$ and $a = 1$): regularity of the spacetime metric requires that we impose a Neumann boundary condition $\partial_a F\big|_{a=0,1} = 0$ on all functions, except for W which has to vanish. To avoid conical singularities, we further require $Y = V$ at $a = 0$ and $Y = U$ at $a = 1$.

- Interpolation boundary (blue dots in Fig. 5.1 (right)): the value of each function is determined by first interpolating the values of the functions in the inner patch, then applying the coordinate transformations.

These boundary conditions are compatible with the DeTurck vector, ξ, vanishing everywhere on the manifold $\mathcal{M}$. However, for stationary spacetimes we do not have a result analogous to that in [48], so we must check *a posteriori* that the components of ξ do indeed vanish to within numerical tolerance.

5.2.4 Technical details

For the data presented in this chapter, the outer patch consists of 320×320 grid points, equispaced in both $0.49 \leq r \leq 1$ and $0 \leq a \leq 1$. The inner patch is constructed by taking a 320×320 grid of points equispaced in both $0 \leq x \leq 1$ and $0 \leq y \leq 1$, then removing points corresponding to $r > 0.51$. We use either fourth- or sixth-order centered difference stencils to discretise the Einstein–DeTurck equations. The resulting non-linear algebraic system is then solved using the Newton line-search method with adaptive step size. We used MUMPS [114, 115] or Intel MKL PARDISO [116, 117] to solve the linear system at each Newton step. The code was based on the nonlinear solver framework (SNES) provided by PETSc [118, 119].

We bootstrap our solution procedure by starting at a high temperature, $\tilde{\kappa} \gg 1$, and a "friendly" angular velocity, $\tilde{\Omega}_H \approx 2$, where we simply use the reference metric as the initial guess to seed the Newton solver. The $\tilde{\kappa} \gg 1$ condition ensures that the effects of the AdS background are small, and so the reference metric, which is built from analytically-known AF solutions, is already almost Einstein in both the near-horizon and near-boundary regions. Meanwhile, the $\tilde{\Omega}_H \approx 2$ condition ensures that the geometry interpolating these two regions is not too highly deformed. These two properties combine to give us the best chance of obtaining convergence from Newton's method. Once we have obtained a solution in this somewhat uninteresting regime, we can use it as the initial guess to seed the Newton solver at less favourable parameters. In this way, we can progressively move away from asymptotic flatness towards the more extreme corners of the parameter space.

5.3 Calculating physical quantities

We now move on to discuss the method that we used to calculate various physical quantities from our solutions. In particular, we chose a rather nontrivial process to calculate the rings' mass, in order to ensure that we obtained an accurate answer.

5.3.1 Horizon area

Our boundary condition manifestly makes the hypersurface $H \coloneqq \{x^a \in \mathcal{M} \mid y = 0\}$ a Killing horizon of $\partial_t - \Omega_H\, \partial_\psi$. To obtain the horizon's area, we simply need to integrate the volume form pulled back onto a constant-t slice of H, thus

$$\begin{aligned} A_H &= \int_{H|_t} \mathrm{d}S\, \sqrt{\det g_{\text{inner}}|_{t,y=0}} \\ &= 4\pi^2 \int_0^1 \mathrm{d}x\, \sqrt{\mathcal{X}_0(x,0)\, \mathcal{R}_0(x,0)\, \mathcal{S}_0(x,0)}\, e^{\frac{1}{2}[\mathcal{X}(x,0)+\mathcal{R}(x,0)+\mathcal{S}(x,0)]}\,. \end{aligned} \tag{5.12}$$

We perform the integration using *Mathematica*'s `NIntegrate` feature, where the unknown functions are first interpolated using polynomial splines.

5.3.2 Angular momentum

Ref. [120] showed that the usual AF Komar integral for angular momentum also gives the correct result in AdS asymptotics. Therefore, we calculate

$$J = \frac{1}{16\pi} \int_\Sigma \star\, \mathrm{d}\big(\partial_\psi\big)^\flat\,, \tag{5.13}$$

where Σ is any closed spacelike 3-surface bounding a region containing the horizon. The full expression for the integrand is complicated and unenlightening so we will not reproduce it here. We choose Σ to reside completely in the outer patch at constant $r \equiv r_J$, where r_J can be any number. In practice, we found that our result varies by less than 1% over the range $0.75 < r_J < 0.95$.

5.3.3 Mass

We found that the most reliable means of calculating the mass is through the first law of black hole mechanics. In terms of nondimensional quantities this reads

$$\mathrm{d}\tilde{M} = \frac{\tilde{\kappa}}{8\pi}\, \mathrm{d}\tilde{A}_H + \tilde{\Omega}_H\, \mathrm{d}\tilde{J}. \tag{5.14}$$

Since each of our data series are obtained by varying $\tilde{\Omega}_H$ while keeping $\tilde{\kappa}$ fixed, we can write the above as an ODE

$$\frac{\mathrm{d}\tilde{M}}{\mathrm{d}\tilde{\Omega}_H} = \frac{\mathrm{d}}{\mathrm{d}\tilde{\Omega}_H}\left(\frac{\tilde{\kappa}\,\tilde{A}_H}{8\pi} + \tilde{\Omega}_H \tilde{J}\right) - \tilde{J}, \tag{5.15}$$

which can be integrated to give

$$\tilde{M}(\tilde{\Omega}_H) = \frac{\tilde{\kappa}\,\tilde{A}_H(\tilde{\Omega}_H)}{8\pi} + \tilde{\Omega}_H \tilde{J}(\tilde{\Omega}_H) - \int_{\tilde{\Omega}_0}^{\tilde{\Omega}_H} \mathrm{d}\omega\, \tilde{J}(\omega) + \tilde{c}_1, \tag{5.16}$$

where $\tilde{\Omega}_0$ is some chosen limit of integration, and $\tilde{c}_1$ is some constant which depends only on $\tilde{\kappa}$. Clearly, the formula (5.16) is not of much use unless we also have a way to fix $\tilde{c}_1$ for each $\tilde{\kappa}$. One way to do this is by integrating (5.14) over a different data series which is continuously connected to an asymptotically flat solution, where we can then use the analytically known mass to fix the constant. To be more precise, we reintroduce explicit factors of the inverse AdS radius $\lambda := \ell^{-1}$ into the first law, thus

$$\mathrm{d}\big(\lambda^2 M\big) = \frac{(\kappa/\lambda)}{8\pi}\,\mathrm{d}\big(\lambda^3 A_H\big) + (\Omega_H/\lambda)\,\mathrm{d}\big(\lambda^3 J\big) \tag{5.17}$$

$$= \mathrm{d}\left[\lambda^2\left(\frac{\kappa\, A_H}{8\pi} + \Omega_H J\right)\right] + \lambda\left(\frac{\kappa\, A_H}{8\pi} + \Omega_H J\right). \tag{5.18}$$

If we now keep κ and Ω fixed while allowing λ to vary, this turns into an ODE which we can integrate with respect to λ

$$M(\lambda) = \left(\frac{\kappa\, A_H(\lambda)}{8\pi} + \Omega_H J(\lambda)\right) + \frac{1}{\lambda^2}\int_0^\lambda \mathrm{d}\lambda'\left[\lambda'\left(\frac{\kappa\, A_H(\lambda')}{8\pi} + \Omega_H J(\lambda')\right)\right] + c_2. \tag{5.19}$$

To fix the constant c_2, note that as $\lambda \to 0$ the integral term becomes

$$\frac{1}{\lambda^2}\int_0^\lambda \mathrm{d}\lambda'\left[\lambda'\left(\frac{\kappa\, A_H}{8\pi} + \Omega_H J\right)_{\lambda=0} + \mathcal{O}(\lambda'^2)\right] \sim \frac{1}{2}\left(\frac{\kappa\, A_H}{8\pi} + \Omega_H J\right)_{\lambda=0}. \tag{5.20}$$

The AF mass is therefore given by

$$M|_{\lambda=0} = \frac{3}{2}\left(\frac{\kappa\, A_H}{8\pi} + \Omega_H J\right)_{\lambda=0}, \tag{5.21}$$

where we have fixed $c_2 = 0$ by identifying the above formula as nothing but the familiar *Smarr relation*. Our mass calculation can thus be summarised as a two-step process:

1. Produce a *"mass calibration"* series of numerical solutions at fixed κ and Ω_H, but over a range of λ. Calculate A and J for each point in the series, then use (5.19) to calculate M. Here, we somewhat arbitrarily chose $\tilde{\kappa} = 0.4$ and $\tilde{\Omega}_H = 0.6$. Note that, in order to calibrate the mass at larger values of λ, we had to increase both κ and Ω_H. This is not a problem, provided that Ω_H/κ remains constant, as we can then divide all three parameters (κ, Ω_H, λ) by the same number to restore constancy. Using this technique, we were able to fix the mass at $\tilde{\Omega}_H/\tilde{\kappa} = 1.5$ for temperatures down to $\tilde{\kappa} = 0.68$.

2. For our *"main"* datasets at some fixed $\tilde{\kappa}$, we can now use (5.16) to calculate the mass by setting $\tilde{\Omega}_0 = 1.5\tilde{\kappa}$ and $\tilde{c}_1 = \left(\tilde{M} - \frac{\tilde{\kappa}\,\tilde{A}_H}{8\pi} - \tilde{\Omega}_H \tilde{J}\right)_{\tilde{\Omega}_H = \tilde{\Omega}_0}$, using the values obtained in step 1.

We close this section by recalling that [121] proved that regular black hole solutions in AdS with a single (positive) angular momentum satisfy a BPS bound:

$$J \leq M\,\ell\,. \tag{5.22}$$

The rotating spherical AdS black holes [106, 43, 107] and perturbative black rings [111] satisfy this bound. Indeed, all of our numerical AdS black ring solutions also satisfy this bound. However, since the black ring becomes singular as $J \to M\,\ell$, the solutions that we have managed to construct never get very close to saturating this limit.

5.4 Geometry

In this section we study the geometry of the spatial cross section of AdS black ring horizons. Throughout, we refer to the size of the rings relative to the AdS radius ℓ, using the terminology defined in §5.1. For small and short black rings, either thin or fat, the horizon geometry is similar to that of the AF black ring. Therefore, we shall not study the horizon geometry of those rings any further. In what follows, we will describe the geometry of black rings which are either long ($R_{S^1} > \ell$) or large ($R_{S^2} > \ell$). In addition, as noted in §5.1, we have not found any evidence that long rings which are both thin and large exist; in other words, our results suggest that all long thin rings are small.

The induced metric on the spatial cross sections of the horizon is

$$\mathrm{d}s_H^2 = R_\parallel(x)^2\,\mathrm{d}\psi^2 + R_\perp(x)^2\,\mathrm{d}\phi^2 + \mathcal{X}_0(x,0)\,e^{\mathcal{X}(x,0)}\,\mathrm{d}x^2\,, \tag{5.23}$$

where $R_\parallel(x) := \sqrt{\mathcal{S}_0(x,0)}\,e^{\frac{1}{2}\mathcal{S}(x,0)}$ and $R_\perp(x) := \sqrt{\mathcal{R}_0(x,0)}\,e^{\frac{1}{2}\mathcal{R}(x,0)}$. To characterise the geometry, it is useful to consider the radii of the rotation circle, R_{S^1}, and of the transverse two-sphere, R_{S^2}. However, unless the ring is very thin, these are rather ambiguously defined, because the S^2 can be highly distorted. It is possible to come up with some reasonable characterisations of these radii that can provide some information about the actual geometry of the horizon. Here we follow [122], and define the inner and outer radii of the horizon S^1 as

$$R_{S^1}^{(\mathrm{inner})} := R_\parallel\big|_{x=1}\,, \qquad R_{S^1}^{(\mathrm{outer})} := R_\parallel\big|_{x=0}\,. \tag{5.24}$$

There are various ways to characterise the size and shape of the S^2. One option is to define the S^2 radius as the radius of its equator, where the S^2 is fattest,

$$R_{S^2}^{(\mathrm{eq})} := \max_{0<x<1}\left(R_\perp(x)\right). \tag{5.25}$$

Alternatively, we can define the radius of the S^2 in terms of the proper length of the S^2 meridians,

$$R_{S^2}^{(\mathrm{mer})} := \frac{1}{\pi}\int_0^1 \mathrm{d}x\,\sqrt{\mathcal{X}_0(x,0)}\,e^{\frac{1}{2}\mathcal{X}(x,0)}\,. \tag{5.26}$$

These two definitions of R_{S^2} coincide when S^2 is perfectly round. For thin rings, these two numbers remain very close, however as the rings become fatter neither of these numbers alone provides an authoritative "size" of the S^2. We can characterise the distortion in the shape of the S^2 by defining the *stretch* σ as

$$\sigma := \frac{R_{S^2}^{(\mathrm{mer})}}{R_{S^2}^{(\mathrm{eq})}} - 1\,. \tag{5.27}$$

A perfectly round S^2 would therefore have $\sigma = 0$. For the black rings, the gravitational self-attraction means that the S^2 is always *prolate*, and so $\sigma \geq 0$. Finally, there is a third definition of the S^2 radius in terms of its area

$$R_{S^2}^{(\mathrm{area})} := \sqrt{\frac{A_{S^2}}{4\pi}}\,, \tag{5.28}$$

where the S^2 area is given by

$$A_{S^2} \coloneqq 2\pi \int_0^1 \mathrm{d}x\, R_\perp(x)\, \sqrt{\mathcal{X}_0(x,0)}\, e^{\frac{1}{2}\mathcal{X}(x,0)}\,. \tag{5.29}$$

In our solutions, we found that these three radii are always related by $R^{(\mathrm{eq})}_{S^2} \leq R^{(\mathrm{area})}_{S^2} \leq R^{(\mathrm{mer})}_{S^2}$.

5.4.1 Isometric embeddings

A useful way to visualise the distortion of the horizon's transverse S^2 is to isometrically embed it into 3D Euclidean space $\mathbb{E}_3$, as was done in [122] (see also [103]). The metric on the horizon S^2 is given by

$$\mathrm{d}s^2_{S^2} = \mathcal{X}_0(x,0)\, e^{\mathcal{X}(x,0)}\, \mathrm{d}x^2 + R_\perp(x)^2\, \mathrm{d}\phi^2\,. \tag{5.30}$$

We wish to embed it into $\mathbb{E}_3$ via cylindrical polar coordinates,

$$\mathrm{d}s^2_{\mathbb{E}_3} = \mathrm{d}u^2 + \mathrm{d}\rho^2 + \rho^2\, \mathrm{d}\phi^2\,. \tag{5.31}$$

Letting

$$u = u(x)\,, \qquad \rho = R_\perp(x)\,, \tag{5.32}$$

the resulting induced geometry is given by

$$\mathrm{d}s^2_{\mathrm{emb}} = \big(R'_\perp(x)^2 + u'(x)^2\big)\, \mathrm{d}x^2 + R_\perp(x)^2\, \mathrm{d}\phi^2\,. \tag{5.33}$$

By comparing (5.33) and (5.30), one obtains the embedding

$$u(x) = \int_0^x \mathrm{d}\eta \sqrt{\mathcal{X}_0(\eta,0)\, e^{\mathcal{X}(\eta,0)} - R'_\perp(\eta)^2}\,, \tag{5.34}$$

which exists for as long as $R'_\perp(\eta)^2 \leq \mathcal{X}_0(\eta,0)\, e^{\mathcal{X}(\eta,0)}$. We find that this condition is satisfied for all the AdS black rings that we have managed to construct.

Figure 5.3 presents embedding plots of some representative AdS black rings. For long thin rings one would expect that the gravitational self-interaction is small, and hence the horizon S^2 should be nearly round. This is precisely what figure 5.3(a) shows for a ring with $\tilde{\kappa} = 5$ and $\tilde{\Omega}_H = 1.3113$.

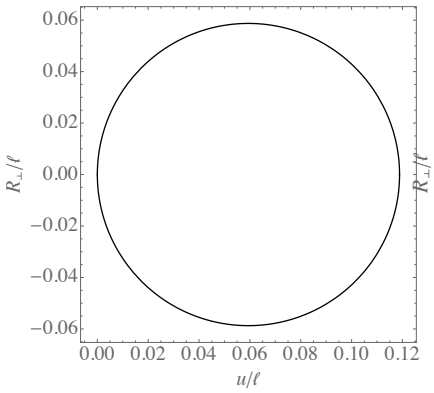

(a) $\kappa\,\ell = 5$, $\Omega_H\,\ell = 1.3113$

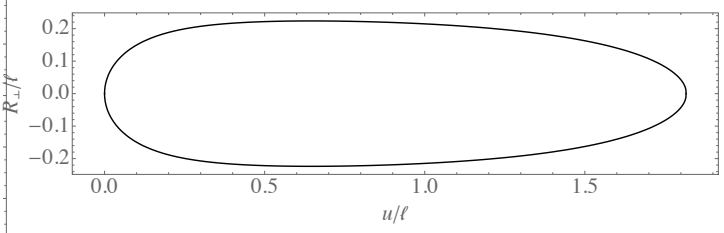

(b) $\kappa\,\ell = 0.5$, $\Omega_H\,\ell = 1.0375$

Figure 5.3: Isometric embeddings. (a) Embedding of a thin ring with $\tilde{R}^{(\text{inner})}_{S^1} = 1.573$ and $\tilde{R}^{(\text{outer})}_{S^1} = 2.097$. The stretch (5.27) is given by $\sigma = 0.006$ and hence the horizon S^2 is almost perfectly round. (b) Embedding of a fat ring with $\tilde{R}^{(\text{inner})}_{S^1} = 0.285$ and $\tilde{R}^{(\text{outer})}_{S^1} = 53.040$. The stretch is $\sigma = 1.812$ and hence the horizon S^2 is highly deformed from spherical symmetry.

Note that the gravitational pull due to the negative cosmological constant is compensated for by having a large enough angular momentum, and hence it should not affect the geometry of the horizon in a significant manner.

Fat rings in AdS have a more interesting geometry. Whilst $R^{(\text{inner})}_{S^1}$ may be small, $R^{(\text{outer})}_{S^1}$ can be very large, and so in this sense it can be *long*. In Fig. 5.3(b) we depict the S^2 embedding of a ring which fits this description: at $\tilde{\kappa} = 0.5$ and $\tilde{\Omega}_H = 1.0375$, we have a rather long $\tilde{R}^{(\text{outer})}_{S^1} = 53.040$. On the other hand, this ring is *not* large in the sense that the typical size of the S^2 is not larger than the radius of AdS. For this particular example we have $\tilde{R}^{(\text{eq})}_{S^2} = 0.224$, $\tilde{R}^{(\text{area})}_{S^2} = 0.415$ and $\tilde{R}^{(\text{mer})}_{S^2} = 0.629$, so the S^2 is indeed highly distorted, as Figure 5.3(b) shows. Note that for any measure of the size of the S^2, we have that $\tilde{R}^{(\text{outer})}_{S^1} \gg \tilde{R}_{S^2}$, so this ring actually looks like a very large and thin membrane with a tiny hole drilled through the middle. It seems reasonable to expect that, by lowering the temperature even further, one should be able to obtain long and fat rings which are also large. But at least in our set up, these are hard to construct numerically.

Figure 5.4 depicts the stretch σ, as defined in (5.27), as a function of the angular velocity $\tilde{\Omega}_H$ for rings at temperatures $\tilde{\kappa} = 0.5, 1., 2.$ (from top to bottom). At sufficiently high temperatures, the $\tilde{\Omega}_H \to 1$ limit is reached by thin rings and hence $\sigma \to 0$ in this limit as the S^2 becomes perfectly

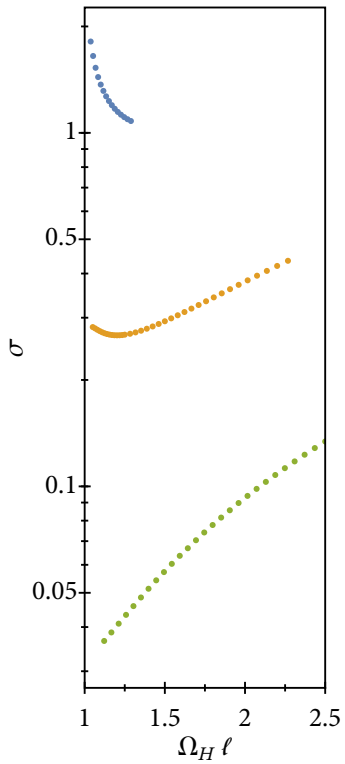

Figure 5.4: The S^2 stretch, σ, plotted against $\tilde{\Omega}_H$, for rings with $\tilde{\kappa} = 0.5, 1, 2$ (top to bottom). At high temperatures, the $\tilde{\Omega}_H \to 1$ limit is reached by thin rings and hence $\sigma \to 0$ in this limit. σ increases monotonically as the ring becomes fatter. At low temperatures, the $\tilde{\Omega}_H \to 1$ limit is reached by membrane rings, so σ cannot be a monotonic function of $\tilde{\Omega}_H$.

round. As the ring becomes fatter, increasing $\tilde{\Omega}_H$ while keeping $\tilde{\kappa}$ fixed, the stretching increases monotonically, since the deformation of the S^2 also increases. It seems natural to expect that σ will diverge in the $\tilde{\Omega}_H \to \infty$ limit. On the other hand, at sufficiently low temperatures, the $\tilde{\Omega}_H \to 1$ limit is reached by the membrane rings. For these temperatures, increasing $\tilde{\Omega}_H$ makes the hole in the middle grow, which implies that σ will *decrease* for a while. However, at some point, the ring starts to become fatter again and hence σ increases.

5.4.2 Invariant radii

Since $R_\perp(x)$ and $R_\parallel(x)$ are both geometric invariants, plotting them against each other allows us to directly compare the relative sizes of the two cycles. However, the information about the lengths along the S^2 meridian is lost and it is therefore important to keep in mind that distances along the curve in these diagrams do not have any real meaning.

Before we describe the geometry of AdS black rings, let us recall some facts about the geometry of rotating AdS black holes. In AdS, the rotating spherical black holes have two different singular limits [111]. In 5D and for fixed mass, the angular momentum of the spherical black hole is strictly less than the BPS value, $J_{\max} < M\,\ell$. In the limit $J \to J_{\max}$ for fixed M, the size of the black hole on the plane of rotation remains finite but the total horizon area goes to zero, hence becoming singular. One can see that in this limit the angular velocity of the horizon diverges. In $D \geq 6$ this corresponds to the well-known ultraspinning limit of black holes, and the value of the angular momentum approaches the BPS value. In AdS, it is possible to take another limit [111], even in 5D, in which both the mass M and the angular momentum J diverge whilst their ratio remains finite with $J/(M\,\ell) \to 1$. In this limit, the black hole approaches a rotating black hyperboloid membrane with a horizon topology $\mathbb{H}^2 \times S^1$.

Figure 5.5 depicts some representative plots for rings in different regimes: thin, fat and membrane. We compare the geometry of the ring with that of the rotating AdS black hole with the same temperature and angular velocity. Note that because we have not fixed the total mass, the actual "sizes" of the black ring and the black hole can be quite different in certain limits. Long thin rings are depicted in 5.5(a). As this plot shows, the radius of the S^1 of the ring is quite large compared to the radius of AdS, and in some sense the black ring is close to the boundary. As we shall see in §5.6, this gets imprinted into the stress-energy tensor of the dual CFT. Figure 5.5(b) shows a fat ring. Even though we could not reliably construct fatter rings at this particular temperature, the plot suggests that the black ring and the black hole would merge in the $\tilde{\Omega}_H \to \infty$ limit. At low enough temperatures, as $\tilde{\Omega}_H \to 1$ the black ring should tend to the same rotating hyperbolic membrane as does the spherical black hole. This is shown in Figures 5.5(c) and 5.5(d). In particular, in Figure 5.5(d), it is quite apparent that the black ring and the black hole are tending to the same solution. Note that, since both the ring and the black hole are close to the *same* black membrane, fixing the mass and the angular momentum instead does not produce a significantly different plot.

5.5 Thermodynamics of AdS black holes

We now move on to discuss the thermodynamics of singly-spinning black holes and black rings in AdS. In §5.5.1 we work in the grand canonical ensemble and study the black hole phases at a fixed

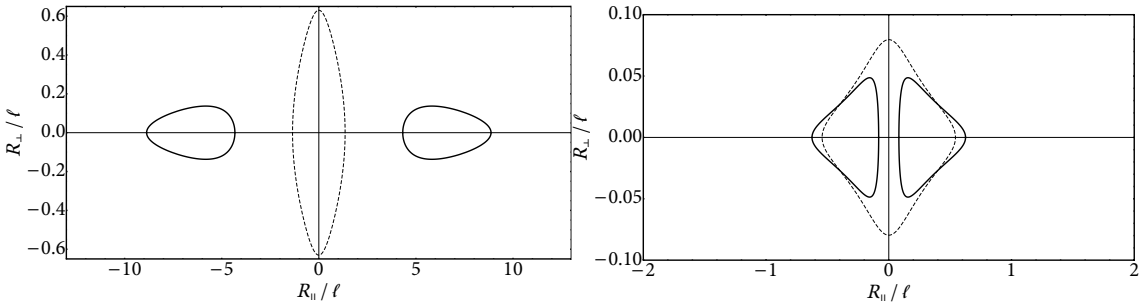

(a) $\kappa\,\ell = 2$, $\Omega_H\,\ell = 1.07704$

(b) $\kappa\,\ell = 2$, $\Omega_H\,\ell = 4.53557$

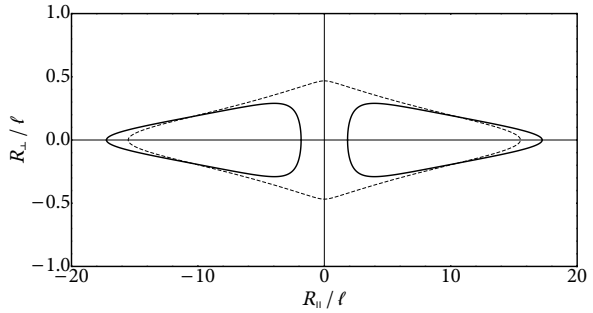

(c) $\kappa\,\ell = 1$, $\Omega_H\,\ell = 1.05048$

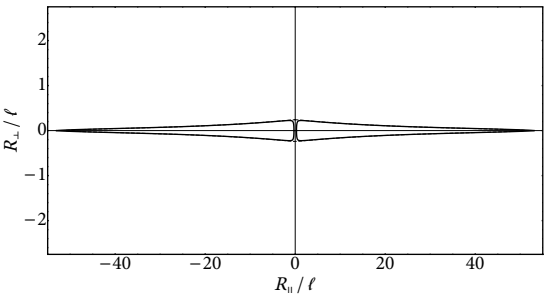

(d) $\kappa\,\ell = 0.5$, $\Omega_H\,\ell = 1.03750$

Figure 5.5: Invariant radii plots for some representative temperatures and angular velocities. The solid line corresponds to the black ring and the black dashed line corresponds to the rotating AdS black hole with the same temperature and angular velocity. In (d) we depict the invariant radii for the same black ring as in Fig. 5.3(b). The aspect ratio is the same in all the plots above, with the vertical axis stretched at exactly 10 times the scale of the horizontal axis.

temperature and angular velocity. Most of our discussion will lie in this section, as this ensemble is much easier for us to access numerically. We also briefly consider the microcanonical ensemble in §5.5.2, where we instead fix the total mass of the solution. This will allow us to make a direct comparison with the perturbative results of [111].

5.5.1 The grand canonical ensemble

The grand canonical ensemble naturally arises from our numerical procedure, as it is precisely the surface gravity and horizon angular velocity that we are able to fix directly as boundary conditions on the horizon. We first review the properties of the rotating spherical AdS black holes, before looking at the thermodynamics of our numerical black ring solutions.

Spherical black holes

We begin by looking at 5D asymptotically AdS solutions for which the metric is known analytically. The most trivial of these is of course the pure global AdS solution itself. This is a solution without a horizon. As a result, we can assign an arbitrary period to the Euclidean time coordinate, and so the temperature can be taken to be anything whatsoever. In this context, it is usually referred to as *thermal AdS*. The other class of solutions that we will consider in this subsection are the topologically spherical black holes rotating in a single plane. These are described analytically by the following metric [43, 107]

$$\begin{aligned} \mathrm{d}s^2 = &-\frac{\Delta}{\rho^2}\left[\mathrm{d}t - \frac{a\sin^2\theta}{1-a^2/\ell^2}\,\mathrm{d}\psi\right]^2 + \frac{\Sigma}{\rho^2}\,\sin^2\theta\left[a\,\mathrm{d}t - \frac{r^2+a^2}{1-a^2/\ell^2}\,\mathrm{d}\psi\right]^2 \\ &+ \frac{\rho^2}{\Delta}\,\mathrm{d}r^2 + \frac{\rho^2}{\Sigma}\,\mathrm{d}\theta^2 + \frac{r^2}{\rho^2}\cos^2\theta\left[r^2\,\Sigma + a^2\cos^2\theta\left(1+\frac{r^2}{\ell^2}\right)\right]\mathrm{d}\phi^2 \\ \rho^2 \coloneqq r^2 + a^2\cos^2\theta \quad &, \quad \Delta \coloneqq \left(r^2+a^2\right)\left(1+\frac{r^2}{\ell^2}\right) - 2\mu \quad , \quad \Sigma \coloneqq 1 - \frac{a^2}{\ell^2}\,\cos^2\theta, \end{aligned} \tag{5.35}$$

where $\mu > 0$ is the mass parameter and a is the rotation parameter. Cosmic censorship requires that $|\tilde{a}| < 1$. The event horizon occurs at $r = r_{\mathrm{H}}$, where r_{H} is the largest real root of the polynomial $\Delta(r)$. After nondimensionalising the parameters,

$$\tilde{\mu} \coloneqq \mu/\ell^2 \quad , \quad \tilde{a} \coloneqq a/\ell \quad , \quad \tilde{r}_H \coloneqq r_H/\ell, \tag{5.36}$$

the physical quantities for these black holes are given by [123]

$$\tilde{M} = \frac{\pi\,\tilde{\mu}\left(3-\tilde{a}^2\right)}{4\left(1-\tilde{a}^2\right)^2}, \tag{5.37}$$

$$\tilde{\kappa} = \tilde{r}_H\left(1+\frac{1+\tilde{r}_H^2}{\tilde{r}_H^2+\tilde{a}^2}\right), \tag{5.38}$$

$$\tilde{A}_H = \frac{2\,\pi^2\,\tilde{r}_H\left(\tilde{r}_H^2+\tilde{a}^2\right)}{1-\tilde{a}^2}, \tag{5.39}$$

$$\tilde{\Omega}_H = \frac{\tilde{a}\left(1+\tilde{r}_H^2\right)}{\tilde{r}_H^2+\tilde{a}^2}, \tag{5.40}$$

$$\tilde{J} = \frac{2\,\tilde{a}\,\tilde{M}}{3-\tilde{a}^2}. \tag{5.41}$$

The solution saturates the BPS bound $|\tilde{J}| \leq \tilde{M}$ as $|\tilde{a}| \to 1$ (although strictly speaking this is a singular limit). It is easy to verify that these quantities do indeed satisfy the first law of black hole mechanics (5.14). From now on we will always take $\tilde{a}$ (and hence $\tilde{\Omega}_H$ and $\tilde{J}$) to be positive.

We begin by solving for $\tilde{a}$ in terms of $\tilde{\kappa}$ and $\tilde{r}_H$:

$$\tilde{a} = \sqrt{\frac{\tilde{r}_H \left(1 - \tilde{\kappa}\, \tilde{r}_H + 2\, \tilde{r}_H^2\right)}{\tilde{\kappa} - \tilde{r}_H}}. \tag{5.42}$$

The BPS limit $\tilde{a} = 1$ corresponds to $\tilde{r}_H = \tilde{r}_H^{(\max)} := \tilde{\kappa}/2$. This is the upper bound on $\tilde{r}_H$ which holds at all temperatures. The static limit $\tilde{\Omega}_H = 0$, i.e. $\tilde{a} = 0$ *and* $\tilde{r}_H \neq 0$, yields two roots $\tilde{r}_H^{(\pm)} := \frac{1}{4}\left(\tilde{\kappa} \pm \sqrt{\tilde{\kappa}^2 - 8}\right)$, corresponding to the *small* and *large* black holes. The singular limit $\tilde{r}_H \to 0$ translates to $\tilde{\Omega}_H \to \infty$.

When $\tilde{\kappa} \geq \sqrt{8}$, both $\tilde{r}_H^{(\pm)}$ are real. In this regime, rotating solutions split into two families: those with $0 < \tilde{r}_H < \tilde{r}_H^{(-)}$ are the *small* rotating black holes, while those with $\tilde{r}_H^{(+)} < \tilde{r}_H < \tilde{r}_H^{(\max)}$ are the *large* rotating black holes. There is *no* solution with $\tilde{r}_H^{(-)} < \tilde{r}_H < \tilde{r}_H^{(+)}$ and so, for a fixed $\tilde{\kappa}$ in this range, the small and large families are not connected.

It is well known that, in the presence of a negative cosmological constant, static black holes cannot exist below the critical *Hawking–Page temparature*, denoted $\tilde{\kappa}_{\rm HP}$ [89]. This is reflected in our calculation here, as the $\tilde{\Omega}_H \to 0$ limit yields imaginary roots when $\tilde{\kappa} < \tilde{\kappa}_{\rm HP} := \sqrt{8} \approx 2.828$. However, rotating black holes *can* still exist at these temperatures as long as they are spinning quickly enough. To see this, we substitute (5.42) into the expression for $\tilde{\Omega}_H$ to obtain

$$\tilde{\Omega}_H^2 = \frac{1}{\tilde{r}_H} \left(\tilde{\kappa} - \tilde{r}_H\right) \left(1 - \tilde{\kappa}\, \tilde{r}_H + 2\, \tilde{r}_H^2\right), \tag{5.43}$$

and hence

$$\left.\frac{\partial \tilde{\Omega}_H^2}{\partial \tilde{r}_H}\right|_{\tilde{\kappa}} = -\frac{1}{\tilde{r}_H^2} \left(\tilde{\kappa} - 3\, \tilde{\kappa}\, \tilde{r}_H^2 + 4\, \tilde{r}_H^3\right). \tag{5.44}$$

The behaviour of $\tilde{\Omega}_H$ as a function of $\tilde{r}_H$ is clearly governed by the cubic factor $C(\tilde{r}_H) := \tilde{\kappa} - 3\, \tilde{\kappa}\, \tilde{r}_H^2 + 4\, \tilde{r}_H^3$, which is always monotonic in $0 \leq \tilde{r}_H \leq \tilde{r}_H^{(\max)}$. While $C(0) = \tilde{\kappa} > 0$ at all temperatures, $C(\tilde{r}_H^{(\max)}) = \frac{1}{4}\, \tilde{\kappa} \left(4 - \tilde{\kappa}^2\right)$ changes sign at $\tilde{\kappa} = 2$. The low temperature solutions are therefore further split into two regimes.

For $2 < \tilde{\kappa} < \tilde{\kappa}_{\rm HP}$ we have $C(\tilde{r}_H^{(\max)}) < 0$, so C must have a root $\tilde{r}_H^*$ corresponding to the turning point in $\tilde{\Omega}_H$. Solutions with $\tilde{r}_H < \tilde{r}_H^*$ are the low-temperature continuation of the *small* black holes family, while those with $\tilde{r}_H > \tilde{r}_H^*$ are the continuation of the *large* black holes family. These two branches *are* now continuously connected to each other at these temperatures.

For $0 < \tilde{\kappa} < 2$ the function $\tilde{\Omega}_H$ has no turning point. Instead, it is monotonically decreasing, with $\tilde{\Omega}_H \to 1$ as $\tilde{r}_H \to \tilde{r}_H^{(\max)}$. Physically, these solutions behave like *small* black holes. Large black holes therefore cease to exist in this temperature regime. It is important to note, however, that the terminology of *small* vs. *large* here no longer reflects the geometric size of the black hole. These "small" black holes can grow to arbitrarily large horizon areas as $\tilde{\Omega}_H \searrow 1$. Instead, we must view it as a classification of thermodynamic stability: *large* black holes are those which are thermodynamically stable.

We can now discuss the phase diagram of these solutions. Fixing the values of $\tilde{\kappa}$ and $\tilde{\Omega}_H$ is analogous to placing the system in the grand canonical ensemble, and thus phase dominance is determined by the grand canonical potential $\tilde{\Phi} := \tilde{M} - \frac{1}{8\pi}\tilde{\kappa}\,\tilde{A}_H - \tilde{\Omega}_H\,\tilde{J}$. With our normalisation, thermal AdS obviously has $\tilde{\Phi} = 0$. The small black holes always have a positive $\tilde{\Phi}$ and therefore never dominate the ensemble. On the other hand, large rotating black holes with negative $\tilde{\Phi}$ do exist at *all* $\tilde{\kappa} > 2$. In the range $2 < \tilde{\kappa} < 3$, some large black holes still have a positive $\tilde{\Phi}$ and so there is an angular velocity threshold below which thermal AdS is still dominant. When $\tilde{\kappa} > 3$, even the static solution has $\tilde{\Phi} < 0$ and so the entire large black hole branch becomes dominant. It is important to note that large rotating black holes always obey the Hawking–Reall bound, $\tilde{\Omega}_H < 1$. This means that they admit a globally defined timelike Killing vector field, thus implying that they should be classically dynamically stable [124]. On the other hand, small black holes violate this bound and could be subject to the superradiant instability, as studied in detail in [125]. The actual phase diagram of rotating AdS black holes is summarised in Figure 5.6.

Black rings

Using the procedure described in §5.3 to calculate physical quantities for the AdS black rings, it is straightforward to compute the associated grand canonical potential. In Figure 5.7, we depict the grand canonical potential for representative AdS black rings at $\tilde{\kappa} = 5$, as a function of the

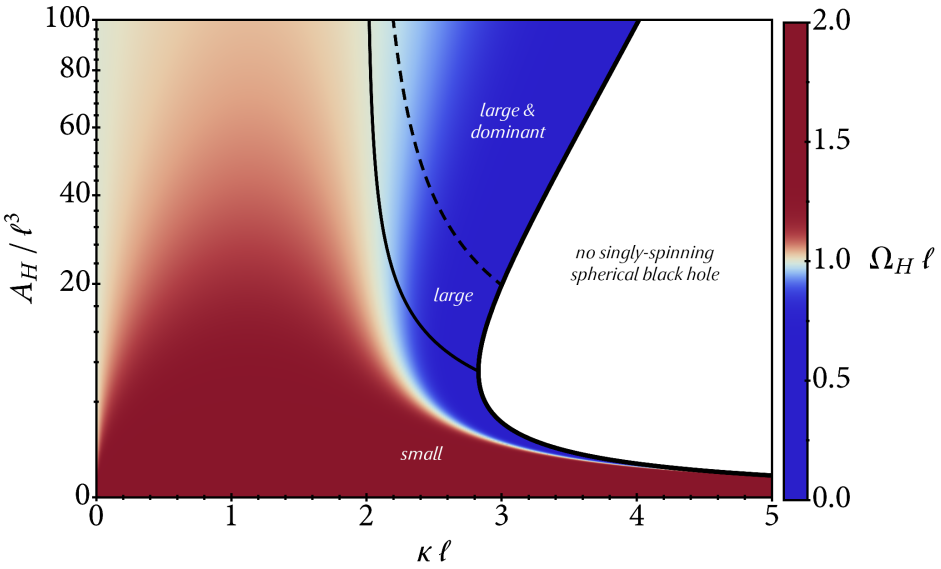

Figure 5.6: A phase diagram for rotating spherical AdS black holes. The colour shows the horizon's angular velocity $\tilde{\Omega}_H$ for each black hole solution, each uniquely parametrised by $(\tilde{\kappa}, \tilde{A}_H)$. Blue-tinted points are *superradiant-stable* solutions with $\tilde{\Omega}_H < 1$. The coloured region is bounded on the right by a thick curve corresponding to the static solutions, which only exist when $\tilde{\kappa} \geq \sqrt{8}$. The solid interior curve marks the boundary between the *small* and *large* black holes, the latter of which cease to exist when $\tilde{\kappa} \leq 2$. Instead, in this regime "small" black holes can have an arbitrarily large $\tilde{A}_H$. Lastly, solutions above the dashed interior curve (all of which are *large* black holes) have negative grand canonical potential $\tilde{\Phi}$, and thus dominate the ensemble.

angular velocity $\tilde{\Omega}_H$, and compare it with that of the *small* rotating AdS black holes at the same temperature. For these configurations, the grand canonical potential of the large black holes is always negative and off the scale of this plot. At *any* other non-zero temperature, the picture for the rings is qualitatively the same. The only difference in the phase diagram for rings as one varies $\tilde{\kappa}$ is that, for $\tilde{\kappa} > 2$, the $\tilde{\Omega} \searrow 1$ limit is attained by thin rings, whilst for $\tilde{\kappa} < 2$ this limit is attained by membrane rings. Therefore, we conclude that black rings in AdS, regardless of their size or shape, *never* dominate the grand canonical ensemble. Moreover, in the $\tilde{\Omega}_H \to \infty$ limit, which is always attained from the fat branch, black rings are connected to the small rotating AdS black holes. Hence, from a thermodynamic point of view, black rings behave in a similar manner to small rotating black holes. In particular, they are always thermodynamically unstable.

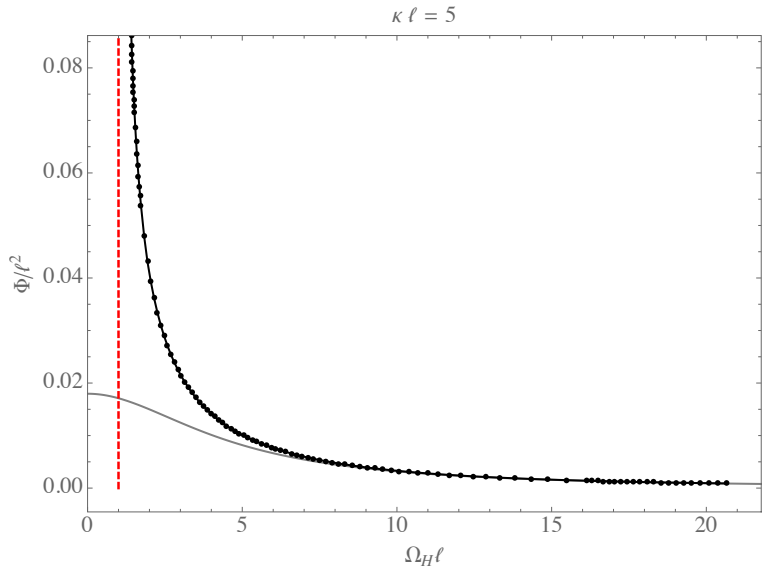

Figure 5.7: Grand canonical potential for AdS black rings (black dots) and small AdS black holes (gray curve) for $\tilde{\kappa} = 5$. The dashed red line indicates the $\tilde{\Omega} = 1$ limit. AdS black rings never dominate the grand canonical ensemble and their angular velocity always exceeds the Hawking–Reall bound.

5.5.2 The microcanonical ensemble

Accessing the microcanonical ensemble is rather more challenging from a numerical point of view. This is because the mass of the black hole is defined in terms of an integral over some hypersurface in the spacetime, which is not a local condition at any given point. We are not aware of any boundary condition which would allow for a specific mass to be fixed directly. For a given mass, we instead resort to first estimating the combinations of the parameters $(\tilde{\kappa}, \tilde{\Omega}_H)$ by interpolating the $\tilde{M}$ across the datasets produced for the grand canonical ensemble. Having obtained solutions near the desired mass at various values of $\tilde{\kappa}$, we then proceed to fine-tune our estimation by performing a bisection search on $\tilde{\Omega}_H$ until the mass becomes correct to within some tolerance. This entails a significant amount of work for each data point (on average we had to obtain four full solutions for each bisection search), and therefore we only present here a phase diagram for one particular choice of $\tilde{M}$.

We choose to focus on $\tilde{M} = 10$ as this allows us to complete the perturbative picture presented

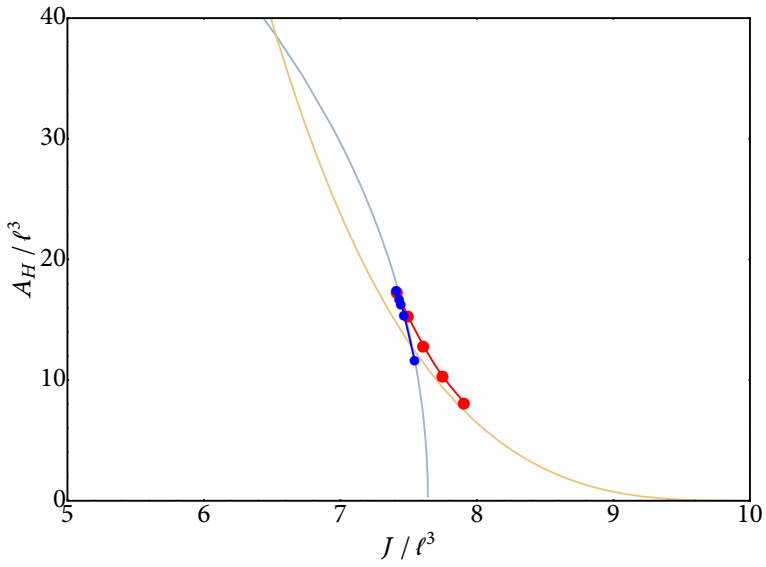

Figure 5.8: Microcanonical ensemble phase diagram for mass $\tilde{M} = 10$, showing the horizon area $\tilde{A}_H$ against angular momentum $\tilde{J}$. The faint blue curve in the background is the exact result for the spherical AdS black holes, while the yellow curve shows the perturbative results of [111]. Red points show our black rings with $\tilde{\kappa} \geq 1$, i.e. the *thin* rings, while the blue points show rings with $\tilde{\kappa} < 1$, i.e. the *fat* rings. There is a cusp at the minimum value of $\tilde{J}$ around $\tilde{\kappa}^* \approx 0.93$. Our numerical results approach the perturbative curve at high $\tilde{\kappa}$, while for low $\tilde{\kappa}$ the fat rings approach the curve for spherical black holes.

in Figure 1 of [111]. For this mass, we have obtained ten solutions in the range $0.6 \leq \tilde{\kappa} \leq 2.0$, and plotted their horizon areas $\tilde{A}_H$ against the angular momenta $\tilde{J}$. This can be superposed onto the corresponding curve for the spherical black holes at the same mass, and also the perturbative result of [111]. As expected, our data points approach the perturbative curve at larger values of $\tilde{\kappa}$, for which the ring is geometrically thin. The BPS bound is approached as $\tilde{\kappa} \to \infty$, where we have $\tilde{A}_H \to 0$, $\tilde{J} \to \tilde{M}$, and the ring becomes arbitrarily thin. Similarly to the AF black rings, the $(\tilde{A}_H, \tilde{J})$ curve for the AdS black rings has a cusp separating "*fat*" and "*thin*" rings in the thermodynamical sense. We estimate that for $\tilde{M} = 10$ this occurs at $\tilde{\kappa}^* \approx 0.93$. As $\tilde{\kappa}$ decreases beyond $\tilde{\kappa}^*$, the ring becomes fatter and the curve approaches that of the spherical black hole, before merging at the singular solution at $\tilde{\kappa} = 0$.

The existence of the cusp allows us to precisely separate *thin rings* from *fat rings* in the microcanonical ensemble. We define *fat rings* by requiring

$$\left.\frac{\partial \tilde{A}_H}{\partial \tilde{\kappa}}\right|_{\tilde{M}} > 0\,. \tag{5.45}$$

One can easily apply the first law (5.14) and the chain rule to deduce that this is equivalent to

$$\Delta := \left.\frac{\partial \tilde{A}_H}{\partial \tilde{\Omega}}\right|_{\tilde{\kappa}} \left.\frac{\partial \tilde{J}}{\partial \tilde{\kappa}}\right|_{\tilde{\Omega}} - \left.\frac{\partial \tilde{A}_H}{\partial \tilde{\kappa}}\right|_{\tilde{\Omega}} \left.\frac{\partial \tilde{J}}{\partial \tilde{\Omega}}\right|_{\tilde{\kappa}} > 0\,. \tag{5.46}$$

It was shown in [126] that the above quantity is proportional to the Hessian determinant of the grand canonical potential $\tilde{\Phi}$ with respect to the coordinates $(\tilde{\kappa}, \tilde{\Omega})$. Note also that Δ carries the same sign as the constant-J heat capacity, C_J.

5.6 Holographic stress tensor

In this final section, we study the stress tensor of $\mathcal{N} = 4$ SYM on $\mathbb{R}_t \times S^3$ for states which are dual to rotating AdS black holes and black rings. For rotating spherical black holes in AdS this was first done in [105], and we shall borrow some results from this reference.

We extract the stress tensor using the standard holographic renormalisation prescription [40]. Note that with our choice of outer patch ansatz (5.5) the boundary geometry is given by the standard metric on the Einstein static universe, $\mathbb{R}_t \times S^3$. As is well known, $\mathcal{N} = 4$ SYM on this geometry has a non-zero Casimir contribution [91]. In the derivation below we will need to subtract this universal piece.

In order to extract the stress tensor of the dual CFT, we must first transform the outer region metric $\mathrm{d}s^2_{\text{outer}}$ from (5.5) into Fefferman–Graham (FG) coordinates. In these coordinates, $\mathrm{d}s^2_{\text{outer}}$ can be expanded around $z = 0$ in the form

$$\mathrm{d}s^2_{\text{outer}} \sim \frac{\ell^2}{z^2}\left(\mathrm{d}z^2 + \left[g_{(0)}(x) + z^2\, g_{(2)}(x) + z^4\, g_{(4)}(x) + \mathcal{O}\big(z^5\big)\right]_{ij} \mathrm{d}x^i\, \mathrm{d}x^j\right), \tag{5.47}$$

$$g_{(0)}(x) + z^2\, g_{(2)}(x) := \mathrm{diag}\Big[-1 - \tfrac{z^2}{2\ell^2},\, \ell^2 - \tfrac{z^2}{2},\, \cos^2\theta\left(\ell^2 - \tfrac{z^2}{2}\right),\, \sin^2\theta\left(\ell^2 - \tfrac{z^2}{2}\right)\Big],$$

$$x^i := (t, \theta, \phi, \psi)\,.$$

Note that index contractions are with respect to the metric $g_{(0)}$, which in this case is just the standard metric on the Einstein static universe. Then, the v.e.v. of the dual stress tensor is given

by [40]

$$\langle T_{ij} \rangle = \frac{\ell^3}{4\,\pi\,G_5} \left(g_{(4)ij} - \frac{1}{8} \left[\left(\operatorname{Tr} g_{(2)} \right)^2 - \operatorname{Tr} g_{(2)}^2 \right] g_{(0)ij} - \frac{1}{2} \left(g_{(2)}^2 \right)_{ij} + \frac{1}{4} \left[\operatorname{Tr} g_{(2)} \right] g_{(2)ij} \right) . \tag{5.48}$$

We perform the change into FG coordinates by substituting $r \to r(z,\theta)$ and $a \to a(z,\theta)$ into (5.5), then imposing the Einstein–DeTurck equations near $z = 0$ at each order in z up to (and including) $\mathcal{O}(z^5)$. We also subtract off the contribution from pure global AdS, i.e. the Casimir energy. The final result, expressed in terms of our unknown functions, is given by

$$\begin{aligned} \left\langle T_{ij} \right\rangle^{\text{sub}} \mathrm{d}x^i\,\mathrm{d}x^j = \frac{N_c^2\,k^8}{768\,\pi^2\,\ell^8} \Big[& - T^{(4)}(\theta)\,\mathrm{d}t^2 - 2\,\ell\,\Xi(\theta)\,\mathrm{d}t\,\mathrm{d}\psi \\ & + \ell^2 \left(Y^{(4)}(\theta)\,\mathrm{d}\theta^2 + \cos^2\theta\,U^{(4)}(\theta)\,\mathrm{d}\phi^2 + \sin^2\theta\,V^{(4)}(\theta)\,\mathrm{d}\psi^2 \right) \Big] , \end{aligned} \tag{5.49}$$

where $F^{(4)}(\theta) := \partial_r^4 F\big(r, \frac{2\,\theta}{\pi}\big)|_{r=k}$ for any function F, and

$$\begin{gathered} \Xi(\theta) := \frac{3\,\pi^4\,\Gamma\left(\Gamma + \tilde{\Omega}_H\right)^2 \left[\left(\Gamma + \tilde{\Omega}_H\right) - \left(\Gamma - 3\,\tilde{\Omega}_H\right) \cos 2\theta \right]^2 \sin^2\theta}{8\,k^4 \left(\Gamma - \tilde{\Omega}_H\right)^4} \left(1 + Z\big(k, \tfrac{2\,\theta}{\pi}\big) \right) , \\ \text{with} \quad \Gamma := \sqrt{4\tilde{\kappa}^2 + \tilde{\Omega}_H^2} \, . \end{gathered} \tag{5.50}$$

Note that our differentiation variable r has dimensions of length, and we have reinstated the compactification scale k, as defined in (5.6), so that the expression above manifestly has the correct dimensions. The expression has already been somewhat simplified by using the relation $T^{(4)}(\theta) + Y^{(4)}(\theta) + U^{(4)}(\theta) + V^{(4)}(\theta) = 0$, which arises from Einstein's equation at $\mathcal{O}(z^2)$. However, we have *not* completely eliminated any of the four functions altogether, so $\left\langle T_{ij} \right\rangle^{\text{sub}}$ is not manifestly traceless. This will prove useful for our calculation method as detailed in §5.6.1. From now on all our stress tensors will have the Casimir contribution subtracted, and as such we no longer explicitly show the "sub" superscripts.

5.6.1 Extraction of the stress tensor

Having derived the stress tensor components in terms of our unknown functions, we now explain our method of evaluating (5.49) numerically. For each unknown function F, we first apply the following protocol. At each $a \in [0,1]$ in the outer patch grid, we take the eight data points

closest to the boundary at $r = 1$ and fit onto them a polynomial of the form $p_a(r) \coloneqq \frac{1}{4!}\,\alpha_a\,(1-r)^4 + \frac{1}{5!}\,\beta_a\,(1-r)^5 + \frac{1}{6!}\,\gamma_a\,(1-r)^6$. The coefficients $\{\alpha_a, \beta_a, \gamma_a\}$ are determined by least squares regression. This fitting naturally has large numerical errors, and so the set $A \coloneqq \{(a, \alpha_a) \mid a \in \text{grid}\}$ must be regarded as a noisy sampling of the fourth derivative $F^{(4)}\big(\frac{\pi\, a}{2}\big)$ evaluated on the boundary $r = 1$. Rather than applying standard noise-reduction filters (e.g. moving averages) on A, we can achieve significantly better results if we take into account the fact that $F^{(4)}$ is a smooth function of θ in the continuum limit. Since we have $\partial_a F^{(4)}\big(\frac{\pi\, a}{2}\big) = 0$ at both $a = 0$ and $a = 1$, we can expand it spectrally as

$$F^{(4)}\big(\tfrac{\pi\, a}{2}\big) = \sum_{n=0}^{\infty} f_n \cos(n\,\pi\, a)\,. \tag{5.51}$$

We can therefore obtain a good approximation to $F^{(4)}$ by fitting the coefficients f_n to the first N terms in the series above. Once again, we determine these f_n by applying least squares regression on the set A. For the work presented here, we managed to achieve good results at $N = 20$.

Next, we note that Einstein's equations imply that we should have

$$\varepsilon(a) \coloneqq T^{(4)}(\tfrac{\pi\, a}{2}) + Y^{(4)}(\tfrac{\pi\, a}{2}) + U^{(4)}(\tfrac{\pi\, a}{2}) + V^{(4)}(\tfrac{\pi\, a}{2}) = 0 \tag{5.52}$$

on the $r = 1$ boundary. Numerical errors mean that we can never expect the functions obtained by fitting f_n as described to yield $\varepsilon(a) \equiv 0$ exactly. However, we noticed that in many cases $\varepsilon(a)$ is actually of the same order of magnitude as the $F^{(4)}$ themselves, even though the DeTurck vector norm $\sqrt{\xi^i \xi_i}$ suggests that these solutions should have very small errors. The nonzero $\varepsilon(a)$ therefore seems to contain some systematic discrepancy beyond what one would expect from pure numerical errors.

With this in mind, we manually enforce (5.52) by subtracting $\varepsilon(a)/4$ from each of the functions $\Big\{T^{(4)}, Y^{(4)}, U^{(4)}, V^{(4)}\Big\}$. One way to gauge the accuracy of our procedure is to calculate the total energy v.e.v. from the stress tensor,

$$\langle E \rangle = -4\,\pi^2\,\ell^3 \int_0^{\pi/2} \mathrm{d}\theta \cos\theta\, \sin\theta\, \big\langle T^t{}_t(\theta) \big\rangle\,. \tag{5.53}$$

To our surprise, this seemingly *ad hoc* procedure resulted in energy densities which, when integrated, agree remarkably well with the black hole masses calculated by the first law method as

described in §5.3.3, with differences ranging between 0.0005% and 0.1%. These results were obtainable by following the above procedure exactly, without having to fine-tune it for each particular solution.

We will not analyse this method in greater detail here, however we will make a few comments. The imposition of (5.52) amounts to solving the leading-order term of Einstein's equations on the AdS boundary. At the computational level, there are infinitely many ways to do this. One could add unequal proportions of $\varepsilon(a)$ to each function, or apply some completely different operations altogether. Our choice corresponds to pulling out a conformal factor from the boundary metric and imposing Einstein's equations by only modifying this conformal factor. We note that this bears a striking resemblance to the conformal decomposition widely used elsewhere in numerical relativity, and it would be interesting to see if a formal justification can be found for its use in this context.

5.6.2 Results

In this subsection we present results for the stress tensor of CFT states dual to black rings for some representative configurations. We concentrate on the energy density distribution on the boundary S^3 and compare it to the energy density distribution of the rotating AdS black hole at the same temperature and angular velocity, using the results of [105].

In order to make the correlation between the horizon geometry and the stress tensor of the boundary CFT apparent, we have depicted in Figure 5.9 the energy density distributions for the same black rings and black holes shown in Figure 5.5. For thin long rings (Figure 5.9(a)) the energy density is concentrated on one of the poles of the boundary S^3, whilst it is negligible on the other pole. The reason is that the ring gets very close to the boundary only along one of the axes of symmetry; the other axis goes through the hole of the ring and hence the energy density in that direction is negligible. Therefore, for such configurations, it becomes particularly simple to distinguish states dual to black rings from states dual to spherical black holes. As the black ring becomes fatter (Figures 5.9(b)-5.9(d)) the energy densities of the black hole and the black ring approach each other, and the distinction between the two becomes less obvious. This is to be expected, since these two phases should eventually merge. This is particularly striking near the membrane limit (Figure

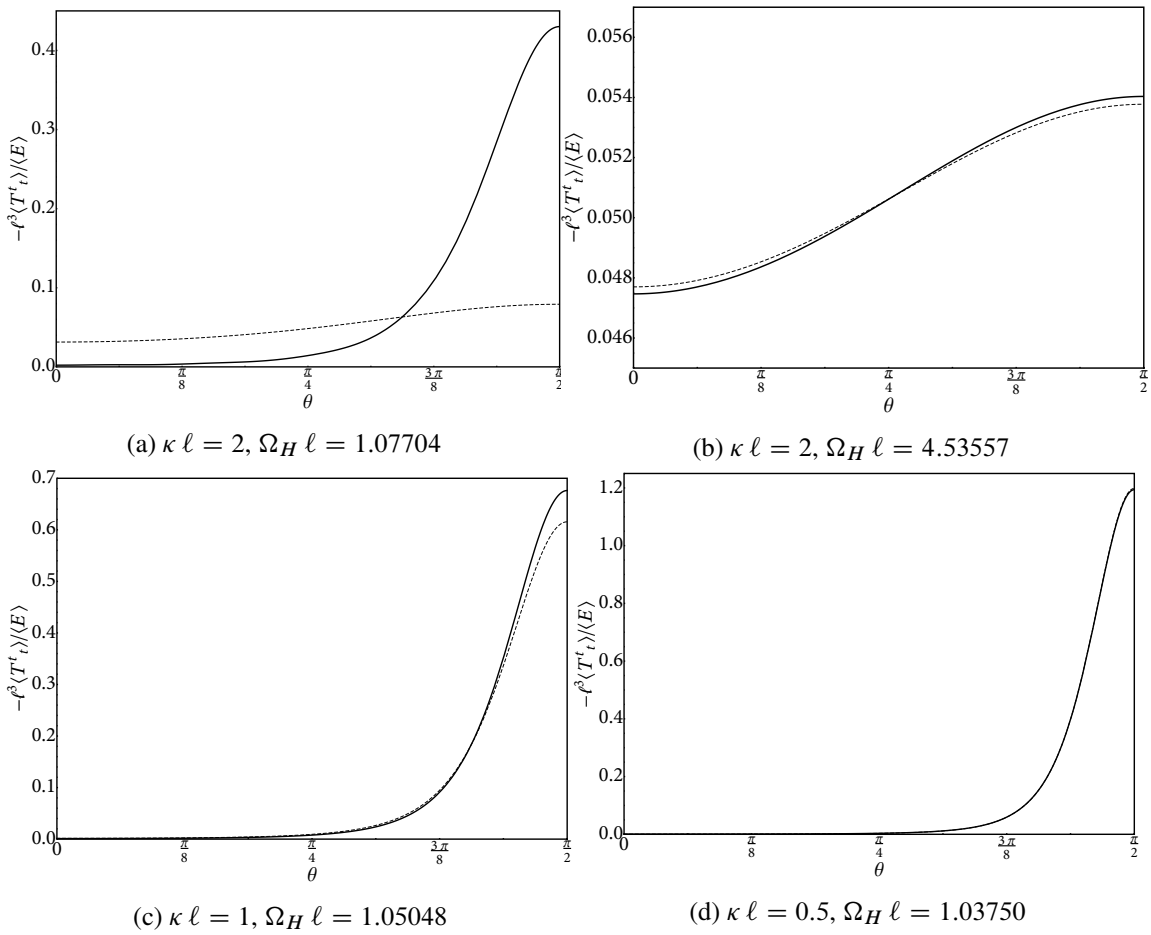

(a) $\kappa\,\ell = 2$, $\Omega_H\,\ell = 1.07704$

(b) $\kappa\,\ell = 2$, $\Omega_H\,\ell = 4.53557$

(c) $\kappa\,\ell = 1$, $\Omega_H\,\ell = 1.05048$

(d) $\kappa\,\ell = 0.5$, $\Omega_H\,\ell = 1.03750$

Figure 5.9: Energy density distribution, normalised by the total energy of the CFT, as a function of the polar angle $\theta = \frac{\pi\,a}{2}$ on the boundary S^3. The data for the black ring is represented by solid black curves, while the data for the corresponding spherical black hole is represented by dashed curves. Firstly, (a) shows the energy density distribution for the same long and thin ring as in Figure 5.5(a). The ring approaches the boundary along the axis of symmetry that goes through one of the poles of the S^2; hence the energy density is concentrated near the corresponding pole of the boundary S^3. By contrast, the spherical black hole exhibits a much more uniform distribution. (c) shows the distribution for a typical fat ring, while (b) and (d) show the distributions for rings approaching the membrane limit. In all of these cases, the energy densities corresponding to the black ring and the rotating black hole are very similar. Indeed, the two curves become virtually indistinguishable as one moves further toward either of these limits.

5.9(d)) for which it is very hard to distinguish the energy density corresponding to the ring from that of the black hole. The fact that we see the energy densities of the two phases approach each other in this limit is reassuring of the correctness of our calculations. Note that in the membrane limit the energy density also gets concentrated on one of the poles. The reason is that in this limit, the bulk solution spreads out on the plane of rotation whilst it becomes infinitely thin in the

transverse directions. Therefore, the energy density should get very large (and eventually diverge) around the pole of the S^3 that connects to the rotation plane, while being negligible around the other pole.

We noted in §5.5 that the thermodynamic behaviour of the AdS black rings is qualitatively similar to that of the small rotating black holes in AdS. This is reflected in the dual stress tensor, in the sense that the latter does not fall into the hydrodynamic regime, even though the stress tensor for rings can be quite different from the stress tensor corresponding to black holes with the same temperature and angular velocity. This result is unsurprising, since [105] solved the relativistic Navier–Stokes equations on S^3 for stationary fluid configurations and found only solutions corresponding to large rotating black holes in AdS.

5.7 Summary

In this chapter, we place AdS black rings among the known stationary black hole phases in the grand canonical ensemble, through the use of numerical techniques. This allows us to cover a large region of the black ring's parameter space, beyond what is accessible perturbatively. The thermodynamic behaviour of AdS rings is qualitatively and quantitatively similar to that of small rotating black holes. The phase diagram for AdS black rings, expressed in terms of T_H and Ω_H, can be summarised as follows:

- $T_H > T_c$: the $|\Omega_H \ell| \to \infty$ limit is reached from the fat family of rings, and the limit is saturated by a singular solution which merges with the small spherical black hole family, as in the AF case. On the other hand, the $|\Omega_H \ell| \to 1$ limit corresponds to an infinitely thin and long ring, and hence is accessible using the perturbation theory of [111].
- $T^* < T_H < T_c$: the limit $|\Omega_H \ell| \to \infty$ still corresponds to the (singular) merger of a fat ring with the spherical black hole. However, the $\Omega_H \ell \to 1$ limit now corresponds to a new membrane-like limit for rings which are *not* geometrically thin. In this limit, black rings tend to the same singular black membrane-type solution as the small black holes, with horizon topology $\mathbb{H}^2 \times S^1$ (see [111]). For intermediate angular velocities, geometrically thin rings can still occur.

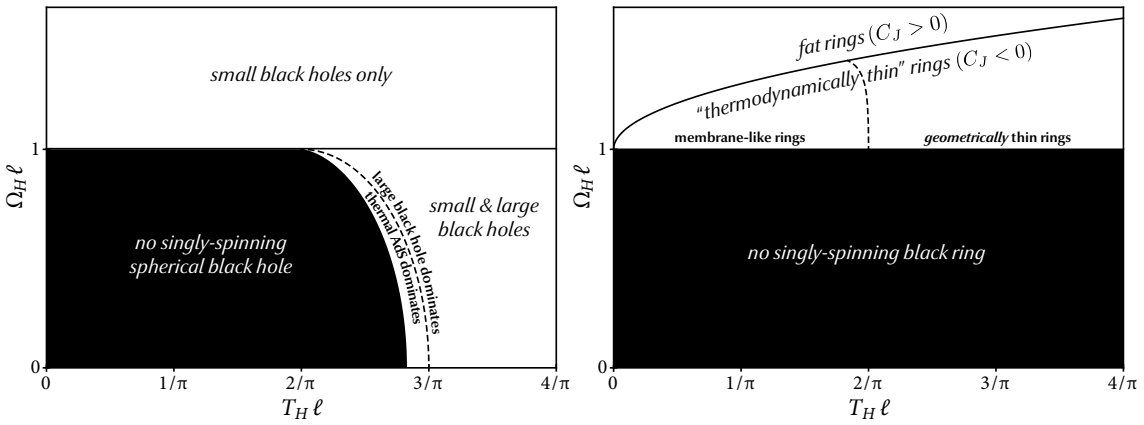

(a) spherical AdS black holes (b) AdS black rings

Figure 5.10: Phase diagrams of singly-spinning (a) spherical black holes and (b) black rings in 5D global AdS, presented in the (T_H, Ω_H) plane. Both kinds of object exist at all non-zero temperatures. All of our AdS black rings have $|\Omega_H\, \ell| > 1$. The black rings can be classified into two distinct phases according to their heat capacity at constant angular momentum, C_J. Those with $C_J > 0$ are *fat rings*, as in the AF case, while those with $C_J < 0$ exhibit somewhat more complicated behaviour. At high temperatures, $T_H\, \ell > 1/\pi$, these rings are indeed *thin*, again as in the AF case. However, for $T_H\, \ell < 1/\pi$, the outer radius of the S^1 of these rings grows much more quickly than its inner radius as $\Omega_H\, \ell \to 1$, causing the S^2 to become highly stretched. In this regime, the ring no longer becomes "thin" in the geometric sense, but instead approaches a membrane-like geometry.

- $T_H < T^*$: there are no geometrically thin rings below a certain temperature T^*. This may be related to the fact that we find no evidence for long thin rings which are also large. The limits $|\Omega_H\, \ell| \to \infty$ and $|\Omega_H\, \ell| \to 1$ are reached by fat and membrane-like rings respectively, as in the previous case. We find that $T^* \ell > \frac{1}{2\,\pi}$.

Figure 5.10 presents a sketch of the various black ring phases on the (T_H, Ω_H) plane. We note that neither fat rings nor membrane-like rings appear to develop new negative modes of the Lichnerowicz operator as one moves along the family of solutions. This suggests that, at least in 5D, these black rings are not unstable under a Gregory–Laflamme type instability along the rotation plane.

Using standard holographic renormalisation techniques [40], we extracted the one-point function for the stress-energy tensor of the dual $\mathcal{N} = 4$ super Yang-Mills (SYM) field on S^3. This observable gives us some idea of how the dual field theory might be affected by the horizon topology of the black hole in the bulk. We find that long thin rings can get very close to the

boundary of AdS, giving rise to an energy distribution that is highly localised on one of the poles of the boundary S^3. This energy distribution is very different from the one corresponding to the small black hole with the same temperature and angular velocity. Unsurprisingly, fat rings give rise to an energy distribution that is qualitatively and quantitatively very similar to that of the corresponding small black hole. In any case, we found that black rings in AdS *never* dominate the grand canonical ensemble. Since AdS black rings always obey $|\Omega_H \, \ell| > 1$, we expect that they should *all* be classically dynamically unstable under superradiance.

CHAPTER 6

Evolving Higher-Dimensional Black Holes with GRCHOMBO

6.1 Introduction

In four dimensions, black hole uniqueness theorems [8, 9, 10, 11] imply that all stationary objects in General Relativity which are sufficiently compact are mathematically described by the Kerr metric [12]. In astrophysics, these objects include, for example, massive bodies at the centre of galaxies and remnants of core-collapse supernovae. The suitability of the black hole description of astrophysical objects implicitly depends on the assumption that the Kerr solution is stable to gravitational perturbations. Mode analysis [127] and the study of the scalar wave equation on the Kerr background [128, 129, 130], along with evidence from numerical relativity [131, 132], strongly suggests that this is indeed the case, however a rigorous nonlinear stability analysis is so far only achieved for the Minkowski spacetime [133]. Very recently, the linear stability of the Schwarzschild solution has been proven in [134], but the corresponding result for the Kerr solution remains an open, and indeed very difficult, problem.

While the landscape of stationary black holes is much richer in higher dimensions, many of these solutions turn out to be unstable. In their seminal paper [30], Gregory and Laflamme demonstrated that black strings and black branes, which are the simple Cartesian products of a Schwarzschild black hole with an arbitrary number of extra flat dimensions, are linearly unstable to gravitational perturbations. The unstable modes have a minimum wavelength, however, and therefore they do not occur if the extra dimensions are compactified to a sufficiently small size. In other words, only

sufficiently thin black strings and branes are subject to the Gregory–Laflamme (GL) instability. Beyond the fact that these solutions are unstable, the linear analysis does not say anything about the dynamical behaviour or the end state of these instabilities. Nevertheless, the authors argued that since the 5D Schwarzshild-Tangherlini black hole has much higher entropy than a long and thin black string, it is possible that the instability could cause the string to pinch off and give rise to a spherical black hole.

It took many more years before a fully nonlinear study of the 5D black string was successfully carried out in [31] using numerical relativity. The remarkable simulation revealed a self-similar process, in which the perturbed black string becomes increasingly more non-uniform and condenses into spherical black hole 'bulges' connected by pieces of thinner black strings. These thinner strings themselves are also subject to GL instability, forming smaller bulges connected by even thinner strings *ad infinitum*. Over the course of their simulation, the authors observed that the timescale between the formation of subsequent generations of bulges decreases geometrically, and therefore by extrapolating this trend they conclude that it is likely that the black string would completely pinch off within a finite time. This change of horizon topology must occur in a singular fashion, and is therefore taken as evidence of a violation of the weak cosmic censorship conjecture [32] in the 5D asymptotically Kaluza-Klein setting.

In asymptotically flat spacetimes, thin stationary black rings [14] locally resemble boosted black strings. Following this observation, [135] extended Gregory and Laflamme's analysis to show that boosted black branes are also subject to the GL instability. They then constructed an approximate model of the black ring as a boosted black string, and concluded that it too should suffer from the GL instability. However, the approximation is only valid for rings with a very small thickness parameter ν, and the authors suggested that their result should really only be trusted when $\nu \lesssim 0.05$, i.e. when the ring is extremely thin. Nevertheless, this result suggests that the existence of a thin black ring could eventually result in a violation of the weak cosmic censorship conjecture in an asymptotically flat setting. Much more recently, a linearised study of gravitational perturbations of the black ring [136] revealed that the linear GL instability indeed affects all thin rings, and in fact even extends into the fat branch as well.

Topologically spherical Myers-Perry (MP) black holes [13], which are the counterpart of Kerr black holes in higher dimensions, are also known to be unstable. In $D > 5$, the rotation of

these black holes is no longer subject to the Kerr-like extremality bound. Instead, they can enter the *ultraspinning* regime with arbitrarily fast rotations [13]. As the black hole rotates faster, the horizon becomes increasingly flattened and eventually starts to resemble a thin black brane [137]. Once again, this led the authors of [137] to the observation that they too would be subject to the GL instability. Ref. [138] subsequently constructed a family of stationary axisymmetric linear perturbations of the MP solution, signalling an onset of GL instability in these black holes. In addition to the axisymmetric GL instability, the authors of [137] also speculated on thermodynamic grounds that MP black holes may also be unstable to non-axisymmetric perturbations, as they have lower entropy than the corresponding system of two boosted Schwarzschild black holes. This was eventually confirmed in the numerical simulations of [139, 33], where it was shown that highly rotating MP black holes are subject to the *bar mode instability*, which tends to stretch the horizon into an elongated shape. The end state of the bar mode instability on 6D black holes with extremely high spins remains an open problem in numerical relativity. It is possible that the horizon would eventually break up in this case, leading to yet another example of violation of weak cosmic censorship.

Certain aspects of these instabilities can be studied using approximate methods. Linearised perturbation theory has been used to successfully identify the onset of potential instabilities [30, 140, 136], while large D calculations [29, 141, 142, 143] can inform us about their dimensional dependence. An obvious advantage of these techniques is that they give a concise, qualitative description of the physics in their relevant domains of validity, covering a range of parameter values in an economical manner. Nevertheless, numerical relativity remains an essential tool with which we can study the end state and the dynamical evolution of these instabilities, including the way in which they interact with each other at the fully nonlinear level. A common feature amongst these scenarios is the dynamical emergence of features at new length scales throughout the course of the instabilities. Fully flexible adaptive mesh refinement (AMR) therefore becomes a necessity for the feasibility and success of these simulations. The GRCHOMBO code was developed to address the need for AMR, both in our study of black hole instabilities and in other areas, such as critical phenomena in gravitational collapse. This is achieved by implementing Einstein's equations on top of the general-purpose CHOMBO AMR library, allowing us to avoid duplicating the significant software engineering effort needed to implement and parallelise numerous complex algorithms.

In this chapter, we begin with a brief review of the CCZ4 formulation of numerical relativity, which gives us a strongly hyperbolic PDE system. We will then give an overview of the adaptive mesh refinement scheme implemented in GRCHOMBO. We will also discuss various techniques which form part of a successful simulation of higher dimensional black holes, namely the use of artificial dissipation and viscosity to control numerical errors, and the *modified cartoon method*, which is used to reduce the simulation to a 2D or 3D computational domain. The final section concerns the considerations involved in tuning runtime parameters to obtain good performance, and in programming Einstein's equations in such a way as to allow the compiler to generate efficient binary code on modern processors.

6.2 The CCZ4 formulation

For a non-stationary metric, the harmonic Einstein equation (2.9) is strongly hyperbolic, and we can therefore directly formulate it into an initial value problem on the metric components $g_{\mu\nu}$. Indeed, this is the basis of the *generalised harmonic coordinates* evolution method, which was successfully employed in the first breakthrough simulation of a binary black hole merger [3]. The main disadvantage of GHC, however, is that it is not robust to singularities on the computational grid. In order to achieve stable evolution, one has to additionally employ *black hole excision*, that is the manual removal of grid points near the singularity. This is typically done by assuming that the physical singularities are contained inside black hole horizons. One would therefore need to find apparent horizons after each time step in order to define the region to be excised. Furthermore, as black holes move through the domain, previous excised points need to be reintroduced, with values extrapolated from neighbouring points.

In order to sidestep the difficulties associated with the handling of singularities in GHC, a large part of the numerical relativity community instead uses '3+1 decompositions' of the Einstein equation with moving punctures. Since we are concerned with simulating higher dimensional spacetimes, we will generalise this terminology to '$d+1$ decomposition'. In this approach, the metric components are first expressed in ADM form

$$\mathrm{d}s^2 = -\alpha(t,x)^2\,\mathrm{d}t^2 + \gamma_{ij}(t,x)\left(\mathrm{d}x^i + \beta^i(t,x)\,\mathrm{d}t\right)\left(\mathrm{d}x^j + \beta^j(t,x)\mathrm{d}t\right). \tag{6.1}$$

Here $\alpha(t,x)$ is called the *lapse* and $\beta^i(t,x)$ is called the *shift* vector. This foliates the spacetime with spacelike hypersurfaces ('slices') Σ labelled by the time coordinate t, where the metric on Σ is given by γ_{ij}. Let $n^\mu := (1/\alpha, -\beta^i/\alpha)$ be the unit normal vector to Σ, we can then define the projectors

$$\gamma_{\mu\nu} = g_{\mu\nu} + n_\mu n_\nu \quad , \quad \gamma^{\mu\nu} = g^{\mu\nu} + n^\mu n^\nu \quad , \quad P^\mu_\nu = \delta^\mu_\nu + n^\mu n_\nu. \tag{6.2}$$

In order to obtain a system of equations which is first order in time, we formulate it in terms of both the induced metric $\gamma_{\mu\nu}$ on Σ, and its *extrinsic curvature* $K_{\mu\nu}$ defined by

$$K_{\mu\nu} := -\frac{1}{2}\mathcal{L}_n \gamma_{\mu\nu}. \tag{6.3}$$

By expanding out the Lie derivative $\mathcal{L}_{\alpha n}\gamma_{ij}$, we can see that

$$\partial_t \gamma_{ij} = -2\,\alpha\, K_{ij} + \mathcal{L}_\beta \gamma_{ij}, \tag{6.4}$$

and so the extrinsic curvature is related to the canonical momentum of the induced metric from the Hamiltonian point of view. We will also define the trace of the extrinsic curvature $K := \gamma^{\mu\nu}K_{\mu\nu}$. Next, we consider various contractions of the Riemann tensor with either P^μ_ν or n^μ to obtain the *Gauss–Codazzi relations*

$$P^\alpha_\mu P^\beta_\nu P^\gamma_\rho P^\delta_\sigma R_{\alpha\beta\gamma\delta} = {}^{(\Sigma)}R_{\mu\nu\rho\sigma} + K_{\mu\rho}K_{\nu\sigma} - K_{\mu\sigma}K_{\nu\rho} \tag{6.5}$$

$$P^\alpha_\mu P^\beta_\nu P^\gamma_\rho n^\delta R_{\alpha\beta\gamma\delta} = D_\nu K_{\mu\rho} - D_\mu K_{\nu\rho} \tag{6.6}$$

$$P^\alpha_\mu P^\gamma_\nu n^\beta n^\delta R_{\alpha\beta\gamma\delta} = \mathcal{L}_n K_{\mu\nu} + K_{\mu\rho}K^\rho{}_\nu + \frac{1}{\alpha}D_\mu D_\nu \alpha, \tag{6.7}$$

where D_μ is the covariant derivative associated with the induced metric on Σ and ${}^{(\Sigma)}R_{\mu\nu\rho\sigma}$ is the Riemann tensor on Σ. In order to apply the Einstein equation, we note also that

$$P^{\mu\rho}P^{\nu\sigma}R_{\mu\nu\rho\sigma} = 2n^\mu n^\nu G_{\mu\nu} \tag{6.8}$$

$$P^{\alpha\mu}n^\nu R_{\mu\nu} = P^{\alpha\mu}n^\nu G_{\mu\nu}, \tag{6.9}$$

where $G_{\mu\nu}$ is the Einstein tensor. By contracting the indices on these relations, equations (6.5)

and (6.6) become the *Hamiltonian constraint* and the *momentum constraint*

$$ {}^{(\Sigma)}R + K^2 - K_{\mu\nu}K^{\mu\nu} = 16\pi\rho \tag{6.10} $$

$$ D_\nu\left(K^{\mu\nu} - \gamma^{\mu\nu}K\right) = 8\pi j^\mu, \tag{6.11} $$

where we define

$$ \rho := n^\mu n^\nu T_{\mu\nu} \tag{6.12} $$

$$ j^\mu := -P^{\mu\nu}n^\rho T_{\nu\rho}. \tag{6.13} $$

These are not evolution equations, but instead impose differential constraints which our functions must satisfy at all time. The remaining equation (6.7) defines the so-called ADM evolution equation. However, this is not yet a numerically stable formulation of the Einstein equation. Indeed, it can be shown that the system is only weakly hyperbolic, leading to an ill-posed initial value problem (see e.g. [144]).

The Baumgarte–Shapiro–Shibata–Nakamura (BSSN) system [145, 146] is a 3+1 formulation of the Einstein equation that is widely popular in the numerical relativity community, which is strongly hyperbolic when combined with appropriate gauge conditions. In BSSN, the induced metric is conformally decomposed and the extrinsic curvature is trace-decomposed. In particular, we define

$$ \begin{aligned} \chi &:= (\det\gamma)^{-1/D} \quad , \quad \tilde{\gamma}_{ij} := \chi\gamma_{ij} \\ A_{ij} &:= K_{ij} - \frac{1}{D}\gamma_{ij}K \quad , \quad \tilde{A}_{ij} := \chi A_{ij} \end{aligned} \tag{6.14} $$

Note that this introduces additional constraints $\det\tilde{\gamma} = 1$ and $\tilde{\gamma}^{ij}A_{ij} = 0$. The variable χ is usually referred to as the *conformal factor*. The crucial modification of BSSN compared to the ADM system is that it treats the conformal connection $\tilde{\Gamma}^i := \tilde{\gamma}^{jk}\tilde{\Gamma}^i_{jk} = -\partial_j\tilde{\gamma}^{ij}$, where $\tilde{\Gamma}^i_{jk}$ is the Christoffel symbol of $\tilde{\gamma}$, as a separate evolution variable. This amounts to adding a multiple of the constraints (6.10) and (6.11) to the ADM equation, which does not modify the physics of the system as these constraints vanish for metrics satisfying the Einstein equation. However, it does modify the character of the PDEs such that they become strongly hyperbolic, leading to a well-posed initial value problem.

BSSN is a *free evolution* scheme, whereby the physical constraints (6.10) and (6.11) are not explicitly imposed at each time step. At the continuum level, it can be shown that if these physical constraints are satisfied in the initial data, then they remain satisfied at all time. However, in numerical solutions, errors mean that constraint violating modes will always be present in the solution. One particular weakness of the BSSN formulation is that there is a Hamiltonian constraint violating mode with a zero speed of propagation [147]. Such a mode would therefore remain on the computational grid throughout the evolution.

From a more mathematical point a view, BSSN modifies the ADM equations in a non-covariant manner. An alternative $d+1$ formulation is proposed in [148], known as the Z4 system, which extends the Einstein equation in a generally covariant way. This formulation is essentially a $d+1$ decomposition of the modified Einstein equation (1.7). Here, the differential constraints (6.10) and (6.11) are replaced by the algebraic constraint $Z_\mu \equiv 0$. Additionally, the components $Z_\mu := (\Theta, Z_i)$, are promoted to become additional evolution variables. Therefore, all constraint violating modes now propagate through the grid. The authors of [149] further showed that the Z4 formulation admits natural constraint damping terms, which can be added covariantly to the Einstein equation via

$$\begin{aligned} R_{ab} + \nabla_a Z_b + \nabla_b Z_a = 8\pi \left(T_{ab} - \frac{1}{D-2} g_{ab} T \right) \\ + \kappa_1 \left(t_a Z_b + t_b Z_a - (1+\kappa_2)\, g_{ab} t_c Z^c \right), \end{aligned} \tag{6.15}$$

where t^a is a timelike vector field. In a $d+1$ decomposition, a natural choice is to take $t^a = n^a$.

The breakthrough success in simulating a dynamical black hole system using BSSN was enabled by the development of the moving puncture gauge and initial data. It is therefore desirable to cast Z4 into a form which would allow us to utilise these standard techniques while still benefiting from its constraint damping properties. This amounts to performing a conformal decomposition of the Z4 system. A covariant way to do this was proposed in [150], resulting in the CCZ4 ("conformal and covariant Z4") system, which is what we implement in GRCHOMBO.[1] Instead of Z^i, the CCZ4 equations are given in terms of the variable $\hat{\Gamma}^i := 2Z^i + \tilde{\Gamma}^i$, where $\tilde{\Gamma}^i$ is the conformal connection

[1] An alternative conformal decomposition, known as Z4c, is given in [151].

as previously defined. In D spatial dimensions, the equations read

$$\partial_t \chi = \frac{2\chi}{D}\left(\alpha K - \partial_k \beta^k\right) + \beta^k \partial_k \chi \tag{6.16}$$

$$\partial_t \tilde{\gamma}_{ij} = -2\alpha \tilde{A}_{ij} + 2\tilde{\gamma}_{k(i}\partial_{j)}\beta^k - \frac{2}{D}\tilde{\gamma}_{ij}\partial_k \beta^k + \beta^k \partial_k \tilde{\gamma}_{ij} \tag{6.17}$$

$$\partial_t K = 2\left(\partial_t \Theta - \beta^k \partial_k \Theta\right) + \alpha\left(\tilde{A}_{ij}\tilde{A}^{ij} + \frac{1}{D}K^2\right) - D^i D_i \alpha + \kappa_1 \alpha\left(1-\kappa_2\right)\Theta \tag{6.18}$$

$$+ 2Z^i \partial_i \alpha + \beta^k \partial_k K \tag{6.19}$$

$$\partial_t \tilde{A}_{ij} = \chi\left[-D_i D_j \alpha + \alpha\left(R_{ij} + 2D_{(i}Z_{j)}\right)\right]^{\mathrm{TF}} + \alpha \tilde{A}_{ij}\left(K - 2\Theta\right) - 2\alpha \tilde{A}_{il}\tilde{A}^l_j \tag{6.20}$$

$$+ 2\tilde{A}_{k(i}\partial_{j)}\beta^k - \frac{2}{D}\tilde{A}_{ij}\partial_k \beta^k + \beta^k \partial_k \tilde{A}_{ij} \tag{6.21}$$

$$\partial_t \Theta = \frac{\alpha}{2}\left[R + 2D_i Z^i - \tilde{A}_{ij}\tilde{A}^{ij} + \frac{D-1}{D}K^2 - 2\Theta K - \kappa_1\left((D+1) + (D-1)\kappa_2\right)\Theta\right] \tag{6.22}$$

$$- Z^i \partial_i \alpha + \beta^k \partial_k \Theta \tag{6.23}$$

$$\partial_t \hat{\Gamma}^i = -2\tilde{A}^{ij}\partial_j \alpha + 2\alpha\left(\tilde{\Gamma}^i_{jk}\tilde{A}^{jk} - \frac{D-1}{D}\tilde{\gamma}^{ij}\partial_j K - \frac{D}{2}\tilde{A}^{ij}\frac{\partial_j \chi}{\chi}\right) + \tilde{\gamma}^{jk}\partial_j \partial_k \beta^i \tag{6.24}$$

$$+ \frac{D-2}{D}\tilde{\gamma}^{ij}\partial_j \partial_k \beta^k + \frac{2}{D}\tilde{\Gamma}^i \partial_k \beta^k - \tilde{\Gamma}^k \partial_k \beta^i + 2\kappa_3\left(\frac{2}{D}\tilde{\gamma}^{ij}Z_k \partial_k \beta^k - \tilde{\gamma}^{jk}Z_j \partial_k \beta^i\right) \tag{6.25}$$

$$+ 2\tilde{\gamma}^{ij}\left(\alpha\partial_j \Theta - \Theta\partial_j \alpha - \frac{2}{D}\alpha K Z_j\right) - 2\alpha\kappa_1 \tilde{\gamma}^{ij}Z_j + \beta^k \partial_k \hat{\Gamma}^i. \tag{6.26}$$

In equation (6.20) above, the superscript TF denotes the trace-free part. The Z4 system and its conformal decompositions are strongly hyperbolic when coupled to the *(1 + log) slicing* and the *gamma-driver shift* gauge conditions

$$\partial_t \alpha = -2\alpha\left(K - 2\Theta\right) + \beta^k \partial_k \alpha \tag{6.27}$$

$$\partial_t \beta^i = F\, B^i + \beta^j \partial_j \beta^i \tag{6.28}$$

$$\partial_t B^i = \partial_t \tilde{\Gamma}^i - \eta B^i + \beta^j \partial_j B^i, \tag{6.29}$$

where B is the time-derivative of the shift, and F and η are some constants.

In order to be able to evolve a black hole without having to explicitly excise grid points in its interior[2],

[2]The author of [152] argued that moving punctures should also be considered a form of "natural excision", where the puncture is implicitly 'excised' by not occurring on any grid point.

we must formulate the initial condition in such a way as to remove physical singularities. For a single black hole, this is typically done by expressing the metric in *(quasi-)isotropic coordinates*, where the initial slice has two asymptotically flat ends, and the black hole's horizon occurs at the throat of the Einstein–Rosen bridge. These coordinates would, however, map the second asymptotic region to a single point (or in the case of the black ring, a circle), thus creating a coordinate singularity there. In a conformally decomposed scheme, this manifests itself as the vanishing of the conformal factor χ. Nevertheless, we can ensure that singularity does not coincide with any grid point, allowing us to avoid evaluating any singular quantity. This forms the basis of the puncture evolution approach.

Another feature of isotropic coordinates is that the resulting lapse function becomes negative inside the horizon, which is not suitable for numerical evolution. A standard practice is to use the *pre-collapsed initial lapse*

$$\alpha_{\text{init}} = \chi^{\frac{D-3}{2}}. \tag{6.30}$$

In the (1 + log) slicing condition (6.27), this causes time to freeze at the black hole's puncture, eventually resulting in the slices becoming highly distorted as constant-coordinate lines fall into the black hole. This phenomenon is known as *slice stretching*. However, the gamma-driver shift condition (6.28-6.29) counters this effect by providing an outward-pointing shift vector at the puncture.

6.3 Adaptive mesh refinement and GRCHOMBO

The presence of multiple, well-separated length scales in the dynamics of black hole instabilities necessitates the use of different grid resolutions in different parts of the computational domain. Covering the entire computational domain with a single, uniformly spaced grid with sufficient resolution for intricate geometric features near the black hole is a prohibitively expensive exercise. Even where it is theoretically possible, it would be an extremely wasteful use of computational resources. Berger–Oliger mesh refinement [153] is by far the most common method used in numerical relativity to achieve multiple grid resolutions. In this method, the computational domain is divided up into a collection of boxes, each containing uniformly spaced grid points. Regions containing finer geometric features would be covered by higher resolution boxes. The resolution

of boxes is refined in discrete levels, with a fixed ratio of grid spacing between the levels. A typical choice is to double the resolution in each level. Instead of *replacing* a box with a higher-resolution one where required, in the Berger-Oliger algorithm we enforce the *proper nesting* property on the boxes: a refined region is covered by boxes in *both* the coarser and finer levels, with the finer-level box entirely contained within the coarser-level one.

The mesh hierarchy can be arranged as either a cell-centered or a node-centered one. A particular mesh refinement library typically only supports one of these modes, or has much better support for one mode over another. The main difference between the two lies in the transfer of data from a finer level to a coarser level. In a cell-centered code, grid points from different levels do not coincide, so we must average the data from all finer grid points within a cell to obtain coarser level data. In a node-centered code, we can simply replace the coarser level data with finer level data at the same grid point. In the opposite direction, data transfer from coarser to finer level is done using polynomial interpolation in both cases. To obtain data at some finer level point, we apply a Lagrange interpolation stencil to the coarser level grid points closest to it. The stencil width is determined by the desired order of convergence.

The organisation of the computational domain into boxes is particularly well-suited to the distributed-memory parallel programming model, as we can readily send different boxes to different compute nodes. However, since we need to apply finite-difference stencils, data from points at the edges of neighbouring boxes must be duplicated into so-called *ghost points*. Ghost data must be re-copied each time the grid data is updated, and since neighbouring boxes can reside in different compute nodes, this represents the biggest serial bottleneck in a mesh-refinement time-evolution code. Ghost data beyond the boundaries of a particular level is obtained by interpolating data from the next coarsest level as described above.

We now provide a high-level overview of the time stepping algorithm. Each level is treated as an independent computational domain when a time step is taken. The numerical relativity community by and large adopts explicit integrators, such as RK4, and therefore we must maintain the CFL convergence condition by using smaller time steps at finer levels. Since our equations are hyperbolic, whenever we reduce the grid spacing by a factor of k, we must also suppress the time step size by a factor of k. Therefore, starting from some time t, if we take a single time step to time $t + \delta t$ at any given level L, then at the next finest level $L + 1$ we must take k steps to reach the

same time. After both levels L and $L+1$ are evolved to time $t+\delta t$, we transfer data from level $L+1$ back to level L in the region where the levels overlap. This procedure is sometimes referred to as "subcycling". An important point to note here is that the recursive nature of this algorithm means that the total number of time steps taken scales exponentially with the number of levels.

In order to take a subcycled time step at a finer level, we first need to fill the out-of-boundary ghost points with interpolated coarser-level data. However, since the coarser level step skips over the subcycled time point, we must also perform an interpolation in time to calculate ghost data. If a Runge-Kutta integrator is used, we can store the intermediate data from the RK substeps in order to achieve a higher order of convergence in the time interpolation process using the so-called *dense output* formula. For RK4, the fourth-order dense output formula between time t and $t+\delta t$ is given by [154, 155]

$$y(t+\theta\,\delta t) = y(t) + \delta t\left[\left(\frac{2\theta^3}{3} - \frac{3\theta^2}{2} + \theta\right)k_1 - \left(\frac{2\theta^3}{3} - \theta^2\right)(k_2+k_3) + \left(\frac{2\theta^3}{3} - \frac{\theta^2}{2}\right)k_4\right], \tag{6.31}$$

which reduces to the standard RK4 step when $\theta = 1$.

Box-structured mesh refinement can be implemented in two different flavours: fixed mesh refinement (FMR) and adaptive mesh refinement (AMR). FMR code requires the user to manually specify the extent of each level beforehand. To simplify this process, the mesh structure generally takes on simple geometric shapes, such as nested cuboids. Once this is done, the refined boxes are allowed to move around to track certain prespecified features, such as a black hole puncture, and they may also be allowed to merge. However, the shape of the refined regions is not altered over the course of the evolution. As the dynamics of the mesh is very simple, FMR codes can be more performant as they have a comparatively smaller overhead, and they also introduce less noise from mesh refinement errors. This approach is therefore very well suited to the simulation of systems which only involve movements of discrete objects whose size do not vary much over time. A prototypical example of this kind of system is the black hole binary. However, in systems whose dynamics involve the creation of new features and length scales, as is the case when simulating instabilities, the FMR approach lacks the flexibility to cope with the ever-changing requirement for the mesh geometry. Moreover, when the object of interest has a non-trivial shape, the manual

specification of the mesh to conform with the geometry or topology of the object can become a painstaking process. On the other hand, if we were to continue to use a geometrically simple mesh hierarchy then we may end up wastefully refining a region which do not require extra resolution. For example, in a simulation of black rings, a hierarchy of nested boxes would lead to the hole in the middle of the ring being unnecessarily refined.

AMR is a more flexible approach which addresses both of these concerns. In AMR, the user specifies a criterion used to determine whether a region requires extra resolution. This is expressed in terms of a tagging function ϕ and a threshold value for ϕ above which a particular point is deemed to be insufficiently resolved. After some user-specified number of time steps, the tagging criterion is reevaluated and a new mesh is created accordingly. Data is then transferred from the old mesh to the new one, and the evolution can continue. The success of an AMR simulation is highly dependent on a good tagging criterion. A criterion which is too conservative would cause the mesh to be refined and unrefined in an erratic manner, injecting large amounts of error into the system in the process, while a criterion which is too eager would lead to a large usage of computational resources. In numerical relativity, it is also important to ensure that the boundary of refined meshes do not lie too closely to the apparent horizon, as this would tend to cause mesh refinement noises to leak into the exterior. The main disadvantage of AMR is clearly the greatly increased complexity in implementation, especially in the context of parallel programming. Fortunately, a number of general-purpose AMR libraries exist on top of which we can build a numerical relativity code with a reasonable level of effort.

Our GRCHOMBO numerical relativity code [156] is built on top of the CHOMBO library [157], which is developed at the Lawrence Berkeley Laboratory. CHOMBO provides a collection of distributed data structures for AMR, along with necessary parallel algorithms implemented through MPI. The library is primarily geared towards cell-centered AMR. The library keeps track of the layout of AMR boxes and handles data transfer both within a node and between nodes, including temporal and spatial interpolation in both directions between coarser and finer levels. The built-in RK4 integrator also distributes the necessary intermediate substeps value across nodes and applies the aforementioned continuity formula to provide a fourth-order accurate temporal interpolation. In order to achieve good performance and scalability, CHOMBO also provides routines to distribute boxes across nodes in a load-balanced manner, however as this can be a fairly costly procedure, the

user has the option to invoke it as often or as rarely as appropriate.

CHOMBO is primarily written in C++, however certain routines which are deemed performance-critical are implemented in Fortran. The user builds a simulation program by subclassing `AMRLevel` and overriding a number of pure-virtual functions to provide the evolution equations, tagging criteria, and to invoke necessary data exchange and interpolation routine. In GRCHOMBO, we implement the CCZ4 equations [150] with Kreiss-Oliger dissipation [158] to reduce high-frequency noises. For the black ring simulation, it was also necessary to add an additional diffusive term near the ring-like singularity in order to stabilise the system. This is done using a localised artificial viscosity term, similar to the ones used in computational fluid dynamics [159]. In order to be able to simulate higher-dimensional spacetimes using a reasonable amount of computational resource, we reduce the effective number of dimensions by exploting rotational symmetries through the modified cartoon method [160, 33]. In the following sections, we give further details of these techniques, along with some technical considerations for efficient implementation and running of the code.

6.4 Artificial dissipation and viscosity

The evolution of a PDE system on a discrete grid with finite resolution can become unstable when the solution contains high-frequency features which cannot be resolved. These features could arise from the nonlinearities in the equations, such as the development of shocks in hydrodynamics, or from numerical errors from domain boundaries. Moreover, when AMR is used, the interpolation between levels could also inject unphysical high-frequency noise into the system. A standard way to address this problem is to modify the equations by adding extra terms to dampen or dissipate modes which cannot be well-resolved by the grid. While such terms would themselves introduce some additional errors into the system, we can nevertheless ensure that they converge to zero as the grid spacing is reduced, at least at the same order as the finite difference scheme used in the evolution. The standard practice in numerical relativity is to add the *Kreiss-Oliger dissipation*

[158] to the right-hand side of each equation,

$$\begin{aligned} \partial_t f &\mapsto \partial_t f + KO \\ KO &:= \frac{\varepsilon}{\delta x} \frac{(-1)^{N-1}}{2^{2N}} D_+^N D_-^N f \end{aligned} \quad , \tag{6.32}$$

where N is some chosen integer, $\varepsilon \leq 1$ is the strength of the dissipation. In one dimension, $D_\pm$ are the one-sided difference operators

$$(D_+ f)_i = -f_i + f_{i+1} \quad , \quad (D_- f)_i = -f_{i-1} + f_i. \tag{6.33}$$

In more than one dimensions, we simply apply these operators along each direction separately. It can be shown that $KO = O(\delta x^{2N-1})$ as $\delta x \to 0$. Therefore, if we use a finite difference scheme in the evolution which converges as $O(\delta x^r)$, then adding the KO term with $N \geq (r+1)/2$ would not change the convergence behaviour. In particular, for our fourth order code we use $N = 3$, where the $D_+^3 D_-^3$ operator can be expressed as a stencil

$$\left(D_+^3 D_-^3 f\right)_i = f_{i-3} - 6f_{i-2} + 15f_{i-1} - 20f_i + 16f_{i+1} - 6f_{i+2} + f_{i+3}. \tag{6.34}$$

Another source of numerical instability in a black hole simulation is the large gradients near the black hole's singularity. In conformally-decomposed evolution schemes such as BSSN or CCZ4, we can aim to contain the singular behaviour entirely within the conformal factor χ, leaving the conformal metric $\tilde{\gamma}_{ij}$ regular. The moving puncture gauge conditions are then employed to move the singularity around the grid, and one typically finds that errors which arise from applying finite difference stencils across the singularity do not leak outside the black hole's horizon. However, where the slicing is not conformally flat, such as in simulations of rotating black holes, the singularity cannot be fully absorbed into χ. While this does not typically cause trouble in simulations of spherical black holes, where the singularity occurs at an isolated point, we found that the extended nature of a black ring's singularity (or, indeed, the disc-like singularity of a Myers–Perry black hole) is much more likely to cause the evolution to become numerically unstable. More specifically, large gradients in the conformal metric components $\tilde{\gamma}_{ij}$ are problematic, as our equations involve the second derivatives of these quantities. In order to achieve stable evolution in this setting, we found it necessary to add an extra diffusive behaviour to $\tilde{\gamma}_{ij}$ on top of the standard

KO dissipation. The form of our diffusion term is inspired by artificial viscosity terms used in computational fluid dynamics, and is given by

$$c_L \, (\delta x)^2 \, g\big(\chi, \left\| \partial \tilde{\gamma}_{ij} \right\|\big) \left(\nabla^2 \tilde{\gamma}_{ij}\right)^{\mathrm{TF}} , \tag{6.35}$$

where c_L sets the strength of the diffusion, and $g\big(\chi, \left\| \partial \tilde{\gamma}_{ij} \right\|\big)$ is some function which is chosen to ensure that the term only becomes important when $\left\| \partial \tilde{\gamma}_{ij} \right\|$ is sufficiently large, and to also confine the effect of this term to a region which is not too close to the horizon. Again, the idea here is to diffuse away modes which cannot be resolved by our computational grid without changing the physical behaviour of the system in the continuum limit. The choice of a traceless diffusion term here ensures that we do not spoil the $\det \tilde{\gamma} = 1$ constraint after each time step. For our simulations, we take

$$g\big(\chi, \left\| \partial \tilde{\gamma}_{ij} \right\|\big) = \left(\mathbb{1}_{\chi \leq \chi_c}\right) \sqrt{\frac{2}{d\,(d-1)} \sum_{i,j,k} \left(\partial_i \tilde{\gamma}_{jk}\right)^2} \quad , \tag{6.36}$$

where χ_c is a threshold for the value of χ above which the diffusion term is deactivated entirely. In our choice of gauge, contours of χ follow roughly the shape of the apparent horizon, and therefore it acts as a good trigger to confine the diffusion region within the horizon's interior. In the 5D black rings simulation, the choice $0.015 < \chi_c < 0.03$ confines the diffused region to well within the horizon, even in the very thin necks when the GL instability is in full swing.

6.5 The modified cartoon method

Numerical computations can quickly become infeasibly costly as the number dimensions increase. At a fixed resolution, the amount of resources required scales exponentially with the number of dimensions. However, many physical behaviours of interest can be studied even when rotational symmetries are imposed. By exploiting these symmetries, we can reduce the effective dimensionality of the problem, which in turn significantly reduces the computational cost involved. One obvious way to proceed is to express the line element in a system of coordinates in which the symmetries are manifest. The downside of this method, however, is the emergence of coordinate singularities at the axes of rotational symmetries, which can become problematic for numerical evolution unless they are carefully handled. Furthermore, the evolution equations will be different

for each component, resulting in a more complex code. Another method, which was implemented in [161], is to perform dimensional reduction on the Einstein equation itself, then perform the $d+1$ decomposition on the resulting modified equation. The *cartoon method* [162] is an alternative approach which allows us to impose symmetries on any $d+1$ decomposition of the Einstein equation while remaining in the relatively simple Cartesian coordinates. This was originally proposed for axisymmetric $d+1$ simulations. The method involves setting up a 2D Cartesian grid to hold the dimensionally reduced data, along with an extra layer of 2D grid above it. Data on these adjacent layers are not evolved using the equations, but are instead filled with values obtained by rotating the actual data grid up, thus imposing the required symmetry. With this extra grid layer, it is possible to calculate derivatives using finite differences as usual. Besides the additional resources required to hold the extra grid, the downside of this approach is that the symmetry is only approximately imposed: the derivatives obtained via finite differences along the reduced dimensions are not what one would get if the symmetries were imposed analytically.

In our work, we employ the so-called *modified Cartoon method* [160, 33], where we continue to work in Cartesian coordinates, but the derivatives along the symmetry-reduced dimensions are now calculated using expressions obtained by imposing the symmetries analytically. We begin by considering a 4-dimensional spacelike slice Σ with Cartesian coordinates (x, y, z, w) and impose a $U(1)$ symmetry on the (z, w) plane. In order to impose the rotational symmetry, we change to polar coordinates on this plane, thus

$$z = \rho\cos\phi \quad , \quad w = \rho\sin\phi, \tag{6.37}$$

where ϕ is the coordinate on the orbits of our $U(1)$ symmetry. Therefore, $\xi := \partial/\partial\phi$ is a Killing vector on Σ. We now choose to perform our computation on the plane $w \equiv 0$, and we denote this subsurface by Σ_0. Our setup is now effectively 3+1, and we use the index $i = 1, 2, 3$ to range over the remaining coordinates (x, y, z). In terms of these Cartesian coordinates, our Killing vector has components

$$\xi = z\,\partial_w - w\,\partial_z. \tag{6.38}$$

The $U(1)$ symmetry can now be expressed as the requirement that the Lie derivatives $\mathcal{L}_\xi$ of all physical quantities on Σ vanish and setting $w = 0$. Expressed in Cartesian coordinates, these conditions can be rearranged to obtain various w-derivatives in terms of grid derivatives on Σ_0.

The resulting relations for the components required for standard CCZ4 can be found in [33], and a complete list all components is given in [163]. For scalar quantities F, we have

$$\begin{aligned}\partial_w F &= 0\\ \partial_i \partial_w F &= 0\\ \partial_w \partial_w F &= \frac{\partial_z F}{z}.\end{aligned}$$

For vector quantities V^a, we have

$$\begin{aligned}
(*) \qquad \partial_w V^i &= -\delta^i_z \frac{V^w}{z} & \partial_w V^w &= \frac{V^z}{z}\\
(*) \qquad \partial_j \partial_w V^i &= -\delta^i_z \left(\frac{\partial_j V^w}{z} - \delta^z_j \frac{V^w}{z^2}\right) & \partial_j \partial_w V^w &= \frac{\partial_j V^z}{z} - \delta_{jz}\frac{V^z}{z^2}\\
\partial_w \partial_w V^i &= \frac{\partial_z V^i}{z} - \delta^i_z \frac{V^z}{z^2} & (*) \qquad \partial_w \partial_w V^w &= \frac{\partial_z V^w}{z} - \frac{V^w}{z^2}.
\end{aligned}$$

For second-rank tensor quantities T_{ab}, we have

$$\begin{aligned}
(*) \qquad \partial_w T_{ij} &= -\delta_{jz}\frac{T_{iw}}{z} - \delta_{iz}\frac{T_{wj}}{z}\\
\partial_w T_{iw} &= \frac{T_{iz}}{z} - \delta_{iz}\frac{T_{ww}}{z}\\
(*) \qquad \partial_w T_{ww} &= \frac{2}{z} T_{zw}\\
(*) \qquad \partial_k \partial_w T_{ij} &= -\frac{1}{z}\left(\delta_{iz}\,\partial_k T_{wj} + \delta_{jz}\,\partial_k T_{iw}\right) + \frac{1}{z^2}\,\delta_{kz}\left(\delta_{iz}\,T_{wj} + \delta_{jz}\,T_{iw}\right)\\
\partial_k \partial_w T_{iw} &= \frac{1}{z}\left(\partial_k T_{iz} - \delta_{iz}\,\partial_k T_{ww}\right) - \frac{1}{z^2}\left(\delta_{kz}\,T_{iz} - \delta_{iz}\,\delta_{kz}\,T_{ww}\right)\\
(*) \qquad \partial_k \partial_w T_{ww} &= \frac{2}{z}\,\partial_k T_{zw} - \frac{2}{z^2}\,\delta_{kz}\,T_{zw}\\
\partial_w \partial_w T_{ij} &= \frac{1}{z}\,\partial_z T_{ij} - \frac{1}{z^2}\left(\delta_{iz}\,T_{zj} + \delta_{jz}\,T_{iz} - 2\,\delta_{iz}\,\delta_{jz}\,T_{ww}\right)\\
(*) \qquad \partial_w \partial_w T_{iw} &= \frac{1}{z}\,\partial_z T_{iw} - \frac{1}{z^2}\left(T_{iw} + 3\,\delta_{iz}\,T_{zw}\right)\\
\partial_w \partial_w T_{ww} &= \frac{1}{z}\,\partial_z T_{ww} + \frac{1}{z^2}\left(T_{zz} - T_{ww}\right)
\end{aligned}$$

More generally, in d spatial dimensions we can take our symmetry group to be $SO(n)$ for any $2 \leq n \leq d$. To fix notation, we take $n = d - 2$ and denote the Cartesian coordinates on Σ by $(x, y, z, w^1, \ldots, w^{n-1})$. As before, the $SO(n)$ acts on the $(z, w^1, \ldots, w^{n-1})$ hyperplane, and our computational domain is the 3-dimensional subsurface Σ_0 obtained by setting all the

w^A coordinates to zero. For $n > 2$, the $SO(n)$ contains $U(1)$ subgroups corresponding to rotations in the (w^A, w^B) planes for any $A \neq B$. The Killing vector generating this subgroup is ${}^{(AB)}\xi := w^A \partial_B - w^B \partial_A$. For any vector V and second rank tensor T, we find that

$$
\begin{aligned}
{}^{(AB)}\mathcal{L}_\xi V^B &= 0 \Rightarrow V^A = 0 \\
{}^{(AB)}\mathcal{L}_\xi T_{iB} &= 0 \Rightarrow T_{iA} = 0 \\
{}^{(AB)}\mathcal{L}_\xi T_{BB} &= 0 \Rightarrow T_{AB} = 0 \\
{}^{(AB)}\mathcal{L}_\xi T_{AB} &= 0 \Rightarrow T_{AA} = T_{BB}
\end{aligned}
$$

Therefore, vector quantities do not carry any extra cartoon components, while tensor quantities only carry a single diagonal cartoon component denoted by T_{ww}. Next, for each w^A, we consider the Killing vector ${}^{(A)}\xi := z\,\partial_A - w^A\,\partial_z$. We now proceed in the same way as in the $U(1)$ case to obtain the same expressions as before, except that those marked with $(*)$ now vanish. Additionally, we find that expressions involving two different cartoon indices vanish, i.e. those with two w indices on the LHS should be replaced by A, B and the RHS multipled by a factor of δ_{AB} or δ^A_B as appropriate. For example, $\partial_j \partial_w V^w \rightarrow \partial_j \partial_B V^A = \delta^A_B \left(\frac{1}{z}\,\partial_j V^z - \frac{1}{z^2}\,\delta_{jz} V^z \right)$. The only expression with a somewhat nontrivial change is the one for $\partial_w \partial_w T_{ww}$, which now reads

$$
\partial_C \partial_D T_{AB} = \frac{1}{z}\,\delta_{AB}\,\delta_{CD}\,\partial_z T_{ww} + \frac{1}{z^2}\,(\delta_{AC}\,\delta_{BD} + \delta_{AD}\,\delta_{BC})\,(T_{zz} - T_{ww})\,.
$$

We can now apply these expressions to the $d+1$ CCZ4 equations to obtain what is essentially the $3+1$ CCZ4 equations with additional terms added involving the cartoon components. From an implementational point of view, this is simply achieved by adding a few extra grid variables for the cartoon components to the code which implements the equations (6.16)-(6.26) (the 'RHS code'). One obvious potential difficulty is the apparent singular behaviour of some cartoon expressions at $z = 0$. While it is possible to regularise the expressions and treat the evaluation at $z = 0$ as a special case, since GRCHOMBO is a cell-centered code, we never actually have grid points at $z = 0$ and therefore we can straightforwardly implement the cartoon expressions.

6.6 Technical considerations

The performance of AMR under distributed-memory parallelism can be highly dependent on the runtime settings. Dividing the domain up into a larger number of smaller boxes means that the computational workload can be straightforwardly distributed to a larger number of nodes, which may appear beneficial in terms of strong scalability of the code. However, more boxes also incur more communication, as we need to fill in ghost data at the boundaries of each box. Furthermore, for very small boxes, the communication overhead can become larger than the computational workload. In GRCHOMBO, we only specify the *maximum* box size, as in certain places the use of large boxes would mean that we cannot conform to the geometry of the solution, and would therefore be refining a larger region than necessary. However, the majority of boxes would still have the maximum box size. For our simulations, we found that setting $32 \leq n_{\max} \leq 64$ allows us to strong-scale well up to 2000 cores.

Another way to increase the performance is to ensure that the RHS code itself is well-optimised. A large part of floating point performance on modern processors relies on *single-instruction, multiple-data* (SIMD) vectorisation. Instead of performing computation on one data point at a time, a SIMD-vectorised code allows the processor to apply the same instruction to multiple data points in each cycle. Compilers are generally capable of turning computational loops into ones which make use of SIMD instructions, however we found that the equations of numerical relativity tend to be sufficiently complicated as to prevent this from happening reliably. As an alternative to automatic vectorisation, we can also program our equations using *compiler intrinsics* which explicitly tell the compiler to generate SIMD instructions, however this approach is tedious and error-prone. The resulting code would also be specific to one particular machine architecture and not portable to any other. Using C++, we can wrap the compiler intrinsics up into overloaded operators, allowing one to explicitly vectorise the code in a much more straightforward manner. To illustrate our solution, let us consider the following simple 2+1 evolution system

$$\partial_t X_{ab} = (X^{-1})^{cd} Y_{ac} Y_{bd},$$
$$\partial_t Y_{ab} = (Y^{-1})^{cd} X_{ac} X_{bd}.$$

This PDE system has no mathematical or physical significance, however it will serve to demonstrate

the issues that arise in the course of writing a numerical relativity code. In the framework of CHOMBO, we need to provide a function which evaluates the RHS of the system. A first attempt at implementing this in C++ could look like the following

```
#define FOR1(i)   for (int i = 0; i < 2; ++i)
#define FOR2(i,j) FOR1(i) FOR1(j)

struct grid_data
{
    double X[2][2];
    double Y[2][2];
}

void invert(const double (&in)[2][2], double (&out)[2][2])
{
    double det = in[0][0]*in[1][1] - in[0][1]*in[1][0];
    out[0][0] =  out[1][1] / det;
    out[0][1] = -out[0][1] / det;
    out[1][0] = -out[1][0] / det;
    out[1][1] =  out[1][1] / det;
}

grid_data calculate_rhs(grid_data in)
{
    double invX[2][2], invY[2][2];
    invert(in.X, invX);
    invert(in.Y, invY);

    grid_data rhs;
    FOR2(a,b)
    {
        FOR2(c,d)
        {
            rhs.X[a][b] += invX[c][d]*in.Y[a][c]*in.Y[b][d];
            rhs.Y[a][b] += invY[c][d]*in.X[a][c]*in.X[b][d];
        }
    }

    return rhs;
}

void rhs_loop(int npoints, const double ** in_arr, double ** rhs_arr)
{
    for (int idx = 0; idx < npoints; ++idx)
    {
        grid_data in = load_data(in, in_arr, idx);
        grid_data rhs = calculate_rhs(in);
        store_data(rhs, rhs_arr, idx);
    }
}
```

We omit the implementation details of the `load_data` function, which loads the input array data at grid point `idx` into a meaningful `grid_data` struct. The `store_data` function does the same

thing in the opposite direction. In the actual code, we also provide functions which apply finite difference stencils to the input array data which also return results as `grid_data` structs. While the code as written above is a reasonably readable and concise representation of our equations, we found that compilers cannot reliably vectorise the outermost loop over the AMR box. Three factors seem to contribute toward this failure, namely the presence of the small inner loops, the use of local structs and arrays, and external function calls. In principle, modern compilers are supposed to be able to handle these situations, however the complexity of the expressions in numerical relativity seems to still confuse the autovectoriser. One obvious solution would be to manually inline the functions, and completely unroll the inner loops and struct declarations, thus

```
void rhs_loop(const double ** in_arr, double ** rhs_arr)
{
    for (int idx = 0; idx < npoints; ++idx)
    {
        double in_X00 = in_arr[0][idx];
        ...
        double in_Y11 = in_arr[7][idx];

        double detX = in_X00*in_X11 - in_X01*in_X10;
        double invX00 = in_X11 / detX;
        ...

        double rhs_X00 =   invX00*in_Y00*in_Y00
                         + 2*invX01*in_Y00*in_Y01
                         + invX11*in_Y01*in_Y01;
        ...

        rhs_arr[0][idx] = rhs_X00;
        ...
        rhs_arr[7][idx] = rhs_Y11;
    }
}
```

In practice, this approach is far too complicated and error-prone for manual coding, and we therefore require a code generator to produce such a function. Mathematica is a good environment in which to create a code generator, thanks to its symbolic manipulation capabilities. The real problem with this approach, however, is that it results in a very large binary when done naively. Instead of structured loops, the compiled code essentially contains multiple copies of the same expression for each component of the tensors. While the code can now be trivially vectorised, it is now significantly larger than the processor's instruction cache, thus causing severe performance degradation. The KRANC code generator [164], which is implemented in Mathematica, follows this type of approach, however it additionally breaks the equation down into self-contained chunks

which can fit in the instruction cache, and can achieve good performance in practice.

An alternative solution is to provide a manually vectorised version of the code using compiler intrinsics. For example, on an AVX machine, we might have

```
struct grid_data_avx
{
    __m256d X[2][2];
    __m256d Y[2][2];
}

void invert_avx(const __m256d (&in)[2][2], __m256d (&out)[2][2])
{
    __m256d det = _mm256_sub_pd(_mm256_mul_pd(in[0][0], in[1][1]),
        _mm256_mul_pd(in[0][1],in[1][0]));

    out[0][0] =  _mm256_div_pd(in[1][1], det);
    ...
}

grid_data_avx calculate_rhs_avx(grid_data_avx in)
{
    __m256d invX[2][2], invY[2][2];
    invert_avx(in.X, invX);
    invert_avx(in.Y, invY);

    grid_data_avx rhs;
    FOR2(a,b)
    {
        FOR2(c,d)
        {
            _mm256_add_pd(rhs.X[a][b], _mm256_mul_pd(invX[c][d],
                _mm256_mul_pd(in.Y[a][c], in.Y[b][d])));
            ...
        }
    }

    return rhs;
}

// Assume for now that npoints is a multiple of 4
void rhs_loop_avx(int npoints, const double ** in_arr, double ** rhs_arr)
{
    for (int idx = 0; idx < npoints; idx += 4)
    {
        grid_data_avx in = load_data_avx(in, in_arr, idx);
        grid_data_avx rhs = calculate_data_avx(in);
        store_data_avx(rhs, rhs_arr, idx);
    }
}
```

In this approach, the code is fully vectorised while still retaining the loop structures of the equations. The resulting binary size is therefore significantly smaller than that of the “fully unrolled” version.

The main disadvantages here are that the code is essentially unreadable and unportable, as we would have to manually create separate versions of the RHS for different hardware platforms. In order to address these issues, we can use object-oriented programming techniques and operator overloading capabilities in C++ to create SIMD-enabled data types which behave like the standard `double`. While we will not reproduce the entire implementation of such a data type here, we will note that care must be taken to ensure that binary operators can accept a mix of SIMD and scalar (both integer and floating point) operands. For example, if `x` is a SIMD vector, then both `2 / x` and `x / 2` should be valid. This can be achieved by declaring appropriate type casts and careful definitions of binary operators in a base type. Our base type is also designed so that the code can fallback to scalar types when no SIMD is detected. A skeletal code for this (we are only implementing division here) is as follows.

```
// Trivial, non-SIMD base template for use
// when no specific SIMD platform is detected
template <typename t>
struct simd
{
    t m_value;

    simd(const t & value) :
        m_value (value)
    {}

    operator t&() { return m_value; }
    static simd load(const double * array) { return *array; }
    static void store(double * array, const simd & v) { _array = v.m_value; }
}

// Helper type to provide platform-specific parameters
// Defaults to scalar unless overridden
template <typename t>
struct simd_traits
{
    using data_t = t;
    static const int simd_len = 1;
}

// Common logic for actual SIMD types
template <typename t>
struct simd_base
{
    using data_t = typename simd_traits<t>::data_t;
    data_t m_value;

    simd_base(const data_t & value) :
        m_value (value)
    {}

    operator data_t&() { return m_value; }
```

```
    // Binary version of the division operator, implemented in terms
    // of the unary version, which is provided as platform-specific code
    friend simd<t> operator/(const simd<t> & a, const simd<t> & b)
    {
        simd<t> out(static_cast<const simd<t>&>(a));
        out /= static_const<const simd<t>&>(b);
        return out;
    }
}
```

Platform-specific logic can then be provided in a separate file as a template specialisation, inheriting the common logic from `simd_base`. An example implementatation for AVX is as follows.

```
template <>
struct simd_traits<double>
{
    using data_t = __m256d;
    static const int simd_len = 4;
}

template <>
struct simd<double> : public simd_base<double>
{
    simd(const double & scalar) :
        simd_base<double> (_mm256_set1_pd(scalar))
    {}

    // C++ does not forward constructor to the base class,
    // so we have to explicitly declare it here
    simd(const __m256d & v) :
        simd_base<double> (v)
    {}

    simd& operator/=(const simd& a)
    {
        m_value = _mm256_div_pd(m_value, a.m_value);
        return *this;
    }

    static simd load(const double * array) { return _mm256_loadu_pd(array); }
    static void store(double * array, const simd & v)
    { _mm256_storeu_pd(array, v.m_value); }
}
```

The above snippet only contains code which is strictly specific to the AVX platform. If we wish to, say, also support SSE packed double, we can simply create another platform-specific source file, and replace the AVX intrinsics with the SSE ones. We can then use compiler macros to detect the target platform at compile time, so that we do not need to worry about selecting the correct platform file manually.

```
#if defined(__AVX__)
#include "avx.hpp"
#elif defined(__SSE2__)
#include "sse.hpp"
#endif
```

Using our SIMD data type, we can reliably generate a fully vectorised program from our original, readable RHS code. The only modification that we need to make is replace `double` with a templated type, and provide a remainder loop for when the array size is not an integer multiple of the SIMD width. For clarity, we use `###` to denote the portion of code which is identical to our first implementation.

```
template <typename data_t>
struct grid_data
{###}

template <typename data_t>
void invert(const data_t (&in)[2][2], data_t (&out)[2][2])
{
    data_t det = ###
    ###
}

template <typename data_t>
grid_data<data_t> calculate_rhs(grid_data<data_t> in)
{
    data_t invX[2][2], invY[2][2];
    ###

    grid_data<data_t> rhs;
    ###
}

void rhs_loop(int npoints, const double ** in_arr, double ** rhs_arr)
{
    int idx = 0;
    int simd_max = simd<double>::simd_len * (npoints / simd<double>::simd_len);

    // SIMD loop
    for (; idx < simd_max; idx += simd_traits<double>::simd_len)
    {
        auto in = load_data<simd<double>>(in, in_arr, idx);
        auto rhs = calculate_rhs(in);
        store_data(rhs, rhs_arr, idx);
    }
```

```
    // Remainder loop
    for (; idx < npoints; ++idx)
    {
        auto in = load_data<double>(in, in_arr, idx);
        auto rhs = calculate_rhs(in);
        store_data(rhs, rhs_arr, idx);
    }
}
```

The driver function `rhs_loop` can now be treated as reusable code. Any threading or optimisation performed on `rhs_loop` could therefore potentially benefit multiple RHS codes. On the other hand, new RHS code can be written in a straightforward manner without the need to consider the compiler's ability to vectorise complex loops.

CHAPTER 7

Finding Apparent Horizons

7.1 Introduction

The ability to locate black hole horizons in a spacetime is a crucial ingredient in any numerical relativity code. This is especially true in our work, where we are interested in the nonlinear dynamics of the black hole horizon itself. An accurate description of the horizon surface is the first step in the study of its physical characteristics. Additionally, from a numerical point of view, this also allows us to ensure that sufficient resolution is used in the black hole's interior. This is critical in achieving a stable evolution and in preventing spurious behaviour from singularity regularisation schemes leaking into the exterior region.

The defining characteristic of a black hole is that it is a region which is causally disconnected from future null infinity. The boundary of this region defines the *event horizon* of the black hole. As this is a global definition, we can only locate the event horizon from the entire history of the spacetime. Generally speaking, this is a prohibitively expensive process, especially in terms of storage and memory requirements. Furthermore, from the numerical perspective, the event horizon plays no role in the local causal structure at each constant time slice Σ, which is what determines the region in which spurious behaviour can be contained. We are instead interested in *apparent horizons*, which are the outermost marginally trapped surface on Σ. More precisely, a *trapped null surface* is a closed, codimension-1 hypersurface $\mathcal{H}$ on Σ, such that the expansion of outgoing null geodesics vanishes everywhere on $\mathcal{H}$, and the apparent horizon is the outermost such surface. The apparent horizon is 'gauge dependent' in the sense that it is specific to a particular choice of

slicing. However, given a particular slice Σ, the apparent horizon is unambiguously defined.

We proceed to present the derivation as given in [144] to obtain a condition satisfied by $\mathcal{H}$. Let s be the outward spacelike unit normal to $\mathcal{H}$ in Σ, and let n be the future-directed timelike unit normal vector to Σ. The outgoing null vector on $\mathcal{H}$ is then given by $l^a := n^a + s^a$. Let $h_{ab} = \gamma_{ab} - s_a s_b$ be the metric pulled back onto $\mathcal{H}$. The expansion of outgoing null geodesics on $\mathcal{H}$ is given by

$$\Theta = \frac{1}{2} h^{ab} \pounds_l h_{ab} = \frac{1}{2} h^{ab} \left(\pounds_s h_{ab} + \pounds_n h_{ab} \right). \tag{7.1}$$

For the first term, we simply note that $-\frac{1}{2}\pounds_s h_{ab}$ is just the extrinsic curvature k_{ab} of $\mathcal{H}$ in Σ. For the second term, we have

$$-\frac{1}{2}\pounds_n h_{ab} = -\frac{1}{2}\pounds_n \gamma_{ab} + \frac{1}{2}\pounds_n (s_a s_b). \tag{7.2}$$

Once again, the first term is just the extrinsic curvature K_{ab} of Σ. Note also that, since $\nabla_a (s_b s^b) \equiv 0$ and $\nabla_a (n_b s^b) \equiv 0$, we also have $h^{ab} \pounds_n (s_a s_b) = 0$. We are therefore left with the condition

$$\Theta = (\gamma^{ab} - s^a s^b)(-k_{ab} - K_{ab}) \equiv 0, \tag{7.3}$$

which the apparent horizon $\mathcal{H}$ must satisfy. In order to proceed, we first need to pick an appropriate parametrisation scheme for $\mathcal{H}$. This turns (7.3) into a PDE system which we can solve numerically. Apparent horizon finders with various different schemes have been implemented, and these are extensively reviewed in [165]. In the following sections, we discuss two different approaches which have been successfully used in studies of higher dimensional black hole instabilities. We begin with the level set approach, which is a very common method in the field, and can be adapted to suit black rings. Limitations of this method will lead us to develop a new, fully parametric approach, which is capable of handling much more general surfaces. As far as we are aware, this method has not been previously used elsewhere.

7.2 Level set parametrisation

A common approach to parametrising apparent horizons in numerical relativity is to find a scalar function F such that $\mathcal{H}$ occurs as the zero contour

$$u - F(v, w) \equiv 0. \tag{7.4}$$

Here (u, v, w) are some coordinates on Σ that are well-adapted to the geometry of $\mathcal{H}$ in its neighbourhood. For example, a popular choice for horizons with spherical topology is the standard spherical polar coordinates $(u, v, w) := (r, \theta, \phi)$, with the radial coordinate r treated as a function of the two angular parameters (θ, ϕ). In this parametrisation, the unit normal to $\mathcal{H}$ is given by $s_a = \nabla_a F / \|\nabla F\|$. By expanding 7.3, we can formulate the apparent horizon equation as a PDE for F

$$\left(\gamma^{ab} - \frac{\nabla^a F \, \nabla^b F}{\|\nabla F\|^2}\right)\left(\frac{\nabla_a \nabla_b F}{\|\nabla F\|} - K_{ab}\right) = 0. \tag{7.5}$$

The PDE can then be solved using Newton's method as described in previous chapters, subject to boundary conditions obtained by requiring that the resulting surface is smooth. This approach can be successfully used to find horizons with a ring-like topology, such as those arising from the evolution of black rings. For this discussion, we denote the grid Cartesian coordinates by (X, Y, Z). We begin by recalling the ring coordinates (x, y, ψ) originally used to describe analytic stationary black rings in [14], which are related to the Cartesian coordinates via

$$\begin{gathered} X = r_2(x, y)\cos\psi \quad , \quad Y = r_2(x, y)\sin\psi \quad , \quad Z = r_1(x, y) \\ r_1(x, y) := \frac{R\sqrt{1 - x^2}}{x - y} \quad , \quad r_2(x, y) := \frac{R\sqrt{y^2 - 1}}{x - y} \\ -1 \leq x \leq 1 \quad , \quad -\infty < y < -1 \quad , \quad 0 \leq \psi < 2\pi, \end{gathered} \tag{7.6}$$

where R is some chosen constant which defines the scale of the coordinate system. More precisely, the coordinate system has a ring singularity at $y = -\infty$, which corresponds to a circle of radius R centered at the Cartesian origin. To gain a geometric understanding of these coordinates, let us fix $\psi = 0$ for the time being and focus on the (X, Z) plane. The flat Euclidean metric is given by

$$\mathrm{d}s^2 = \frac{R^2}{(x - y)^2}\left(\frac{\mathrm{d}x^2}{1 - x^2} + \frac{\mathrm{d}y^2}{y^2 - 1}\right). \tag{7.7}$$

Near the singularity, i.e. as $y \to -\infty$, we have

$$\mathrm{d}s^2 \sim R^2 \left(\frac{\mathrm{d}y^2}{y^4} + \frac{\mathrm{d}x^2}{y^2\left(1-x^2\right)} \right) = R^2 \left(\mathrm{d}(1/y)^2 + (1/y)^2 \, \frac{\mathrm{d}x^2}{1-x^2} \right). \tag{7.8}$$

If we now define new coordinates (ρ, ξ) with $x = \cos\xi$ and $\rho \sim -1/y$ as $y \to \infty$, we would get

$$\mathrm{d}s^2 \sim R^2 \left(\mathrm{d}\rho^2 + \rho^2 \, \mathrm{d}\xi^2 \right). \tag{7.9}$$

In other words, (ρ, ξ) form a system of local polar coordinates about the ring singularity. Furthermore, it is desirable also have $y \to -1$ as $\rho \to \infty$. In our code, we choose $y = -1/\tanh\rho$, however presumably other choices such as $y = -1/\rho - 1$ could also work. The horizon $\mathcal{H}$ is then described as the zero set $\rho - F(\xi, \psi) \equiv 0$, with the function F satisfying the PDE (7.5). The problem with this setup, however, is that it is not suitable when $\mathcal{H}$ becomes highly stretched. If we cannot fit a circle of any size into the interior of the projection of $\mathcal{H}$ onto the (X, Y) plane, then the (ρ, ξ, ψ) coordinates as currently defined cannot be used to parametrise $\mathcal{H}$ under this approach, as the ring singularity must intersect $\mathcal{H}$ regardless of the chosen value of R. One solution to overcome this limitation is to allow R to be a function of ψ rather than a constant. We would then have to inspect the grid functions to guess the shape of $\mathcal{H}$, e.g. by looking at the contour of the

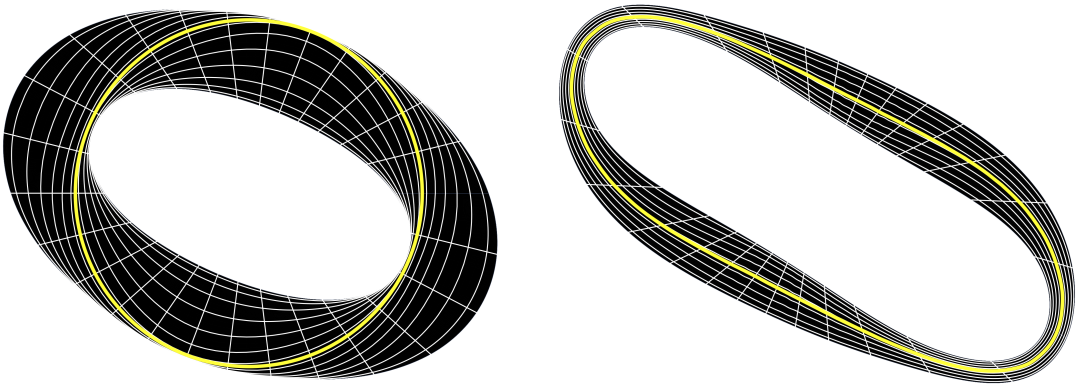

Figure 7.1: Standard ring coordinates with constant R are not suitable when the apparent horizon is highly deformed. The left plot shows the greatest extent of deformation which can be handled by these coordinates. Clearly, if the surface becomes any more eccentric then we can no longer fit the circular coordinate singularity entirely within $\mathcal{H}$. The severe distortion of coordinate lines where the singularity is very close to $\mathcal{H}$ also causes steep gradients which are problematic for the numerical solver. The same coordinate system becomes more generally applicable if R is allowed to be a function of ψ. The right plot shows a much more eccentric $\mathcal{H}$ which can be parametrised by ring coordinates when the function $R(\psi)$ is chosen appropriately. Note that the coordinate lines are also much more uniformly distributed across the surface.

conformal factor χ, and manually specify the function $R(\psi)$ to roughly conform to the expected geometry. To summarise, we can use the following coordinate system to parametrise a fairly large class of ring-like horizons as a level set:

$$\begin{gathered} X = r_2(\rho,\phi,\psi)\cos\psi \quad , \quad Y = r_2(\rho,\phi,\psi)\sin\psi \quad , \quad Z = r_1(\rho,\phi,\psi) \\ r_1(\rho,\phi,\psi) := \frac{R(\psi)\sin\xi}{\cos\xi + 1/\tanh\rho} \quad , \quad r_2(\rho,\phi,\psi) := \frac{R(\psi)/\sinh\rho}{\cos\xi - 1/\tanh\rho} \\ 0 \le \rho < \infty \quad , \quad -\pi \le \phi \le \pi \quad , \quad 0 \le \psi < 2\pi, \end{gathered} \tag{7.10}$$

where $R(\psi) > 0$ is some prescribed function that roughly follows the shape of $\mathcal{H}$. Creating a suitable function $R(\psi)$ for each $\mathcal{H}$ is a rather tedious process if done manually, however it can be somewhat automated by using some heuristic algorithm, e.g. setting R to the point with minimum χ along each constant-ψ ray.

The main limitation of this approach is that it can only describe surfaces which occupy a convex region in the (v, w)-plane. From a practical point of view, we typically further restrict this to be a rectangular region. (For topologically spherical horizons, this translates to the requirement that the projection of $\mathcal{H}$ onto the equatorial plane of the coordinate system forms a *star domain*.) When simulating rapidly spinning black holes, however, frame dragging effects tend to twist $\mathcal{H}$ in such a way as to break the convexity condition. In this case, we can apply some intuition to distort the coordinate system to compensate. The simplest way to do this is to define *twisted Cartesian* coordinates $(\tilde{X}, \tilde{Y}, \tilde{Z})$ with $Z = \tilde{Z}$ and

$$\begin{aligned} X &= \tilde{X}\cos\left(kr^2\right) - \tilde{Y}\sin\left(kr^2\right) \\ Y &= \tilde{X}\sin\left(kr^2\right) + \tilde{Y}\cos\left(kr^2\right) \quad . \\ r^2 &:= X^2 + Y^2 = \tilde{X}^2 + \tilde{Y}^2 \end{aligned} \tag{7.11}$$

By replacing the Cartesian coordinates in (7.10) by the twisted ones, we can continue to use the ring coordinates to parametrise $\mathcal{H}$ which is subject to some mild frame dragging.

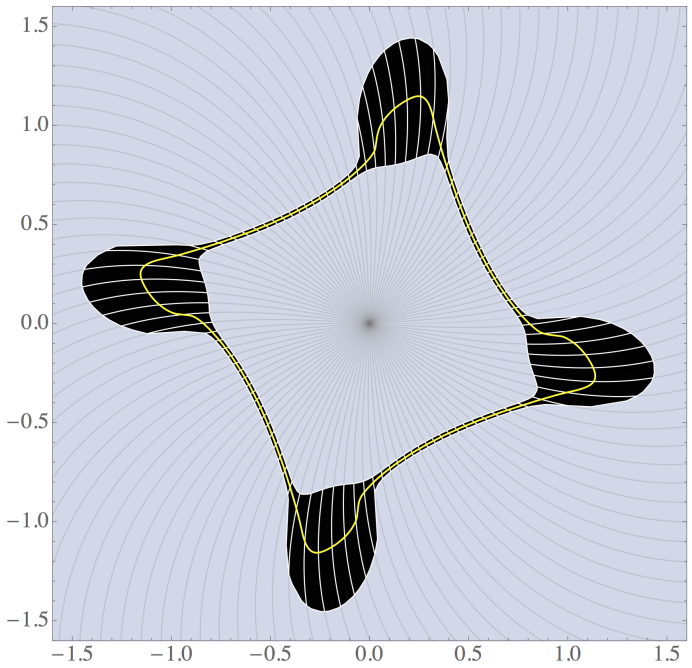

Figure 7.2: By applying a twisting transformation to the angular coordinates, we are able to use the ring coordinates to parametrise a mildly frame-dragged apparent horizon using the level set approach. However, this is only applicable in a very limited setting.

7.3 General parametrisation

The intuition-led coordinate transformation described in the previous section is an *ad hoc* solution to address a major limitation of the level set approach. In the study of black hole instabilities, we may encounter complex geometries for which there is no obvious coordinate trick that applies. Figure 7.3 shows one scenario in which this occurs.

In order to create a more robust apparent horizon finder, we propose to instead treat $\mathcal{H}$ as a general parametric surface. More precisely, we define $\mathcal{H}$ via

$$x^i = X^i(u^\alpha), \tag{7.12}$$

where x^i $(i = 1, \ldots, d)$ are the Cartesian coordinates on the constant-time slice Σ, and the u^α $(\alpha = 1, \ldots, d-1)$ are parameters on the surface $\mathcal{H}$. From a computational point of view, we now need to solve for d unknown functions, X^i, in order to determine $\mathcal{H}$. The tangent and normal

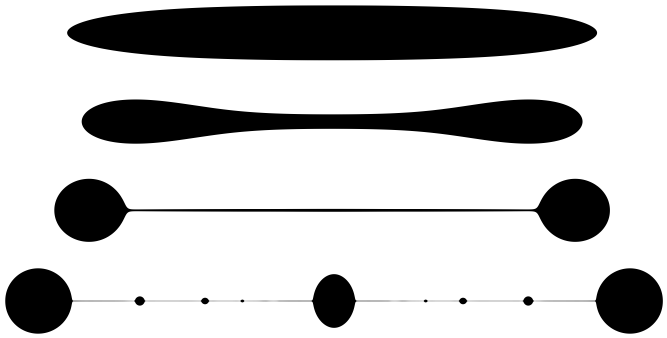

Figure 7.3: In the initial stages of the evolution of an axisymmetrically perturbed, ultraspinning 6D Myers-Perry black hole, the apparent horizon forms a star domain and can therefore be constructed using the level set parametrisation with standard spherical coordinates. This is clearly not the case in the later stages as the GL instability develops in earnest. There is no obvious coordinate system which can describe the dynamically changing geometry as a level set, and we must instead resort to using a more general parametrisation in order to construct the horizon.

vectors to $\mathcal{H}$ in Σ are given by

$$T^i_{(\alpha)} = \frac{\partial X^i}{\partial u^\alpha} \quad \text{and} \quad S^i = \star \left(T_{(1)} \wedge \cdots \wedge T_{(d-1)} \right)^\sharp \tag{7.13}$$

respectively. For notational convenience, we will also use $t^i_{(\alpha)}$ and s^i to denote the corresponding unit vectors. The extrinsic curvature of $\mathcal{H}$ in Σ is given by

$$\begin{aligned} k_\alpha &= -T^i_{(\alpha)} T^j_{(\beta)} \nabla_j s_i \\ &= -T^j_{(\beta)} \nabla_j \left(s_i \, T^i_{(\alpha)} \right) + s_i \, T^j_{(\beta)} \nabla_j T^i_{(\alpha)} \\ &= s_i \left(\frac{\partial^2 X^i}{\partial u^\alpha \partial u^\beta} + \Gamma^i_{jk} \frac{\partial X^j}{\partial u^\alpha} \frac{\partial X^k}{\partial u^\beta} \right). \end{aligned} \tag{7.14}$$

Denoting the induced metric on $\mathcal{H}$ by $\gamma_{\alpha\beta} = T^i_{(\alpha)} T^j_{(\beta)} \gamma_{ij}$ and its inverse $\gamma^{\alpha\beta} = \left(\gamma_{\alpha\beta} \right)^{-1}$, the equation for $\mathcal{H}$ can be expressed as

$$s_i \, \gamma^{\alpha\beta} \left(\frac{\partial^2 X^i}{\partial u^\alpha \partial u^\beta} + \Gamma^i_{jk} \frac{\partial X^j}{\partial u^\alpha} \frac{\partial X^k}{\partial u^\beta} \right) + \left(\gamma^{ij} - s^i s^j \right) K_{ij} = 0, \tag{7.15}$$

where we treat s^i as a function of the $\partial X^i \big/ \partial u^\alpha$.

Since we now need to solve for d functions, we must provide an additional set of $d - 1$ equations in order to complete the PDE system. We begin by considering the one-dimensional case, where

$\mathcal{H}$ can be described as a curve on an (x, z) plane, parametrised by a single parameter u. One way to fix the gauge is to fix the norm of the tangent vector to some prescribed values $\tilde{H}(u)$ along the curve, that is,

$$\gamma_{ij} \frac{\partial X^i}{\partial u} \frac{\partial X^j}{\partial u} = \tilde{H}(u). \tag{7.16}$$

However, imposing the above condition amounts to fixing the coordinate distance between the two ends of the curve, which is not generally known *a priori*. Instead, we differentiate it to form a second order condition

$$\begin{aligned}
\frac{1}{2} \frac{\partial X^k}{\partial u} \nabla_k \left(\gamma_{ij} \frac{\partial X^i}{\partial u} \frac{\partial X^j}{\partial u} \right) &= \gamma_{ij} \frac{\partial X^j}{\partial u} \frac{\partial X^k}{\partial u} \nabla_k \left(\frac{\partial X^i}{\partial u} \right) \\
&= \gamma_{ij} \frac{\partial X^j}{\partial u} \frac{\partial X^k}{\partial u} \left(\frac{\partial u}{\partial x^k} \frac{\partial^2 X^i}{\partial u^2} + \Gamma^i_{kl} \frac{\partial X^l}{\partial u} \right) \\
&= \gamma_{ij} \left(\frac{\partial^2 X^i}{\partial u^2} \frac{\partial X^j}{\partial u} + \Gamma^i_{kl} \frac{\partial X^j}{\partial u} \frac{\partial X^k}{\partial u} \frac{\partial X^l}{\partial u} \right) \\
\Rightarrow \quad \gamma_{ij} \left(\frac{\partial^2 X^i}{\partial u^2} \frac{\partial X^j}{\partial u} + \Gamma^i_{kl} \frac{\partial X^j}{\partial u} \frac{\partial X^k}{\partial u} \frac{\partial X^l}{\partial u} \right) &= H(u),
\end{aligned} \tag{7.17}$$

where $H(u)$ is some prescribed function. A simple choice is to set $H(u) \equiv 0$, which is equivalent to requiring that u be an affine parameter. This condition and (7.15) together form a system of well-posed PDEs which can be solved numerically. Since $\mathcal{H}$ is a compact surface in Σ, the domain of u must either be periodic or terminate on an axis of symmetry. In this latter case, boundary conditions are determined by smoothness requirements, such as the absence of conical singularities.

As usual, we apply the Newton line search method to solve the nonlinear PDE system. In our current setup, both the expansion and the gauge condition are simultaneously corrected in each Newton step. When the geometry of $\mathcal{H}$ becomes more extreme, we found that the gauge condition becomes significantly more stiff than the actual physical equation, and the nonlinear solver either requires strongly suppressed step sizes to make progress in the best case, or outright diverges in the worst case. We found that the Newton solver becomes significantly more robust if we project out the gauge mode from the line search direction entirely. To do this, instead of choosing a particular gauge by providing a specific $H(u)$ in (7.17) *a priori*, we now set $H(u)$ to coincide with the LHS of (7.17) at the current iteration. In other words, we are fixing the gauge to be whatever gauge the

current iteration is in. Figure 7.4 illustrates the difference this makes in a schematic diagram. In practice, this is very easy to implement in code: we simply need to fix the Newton residual of the gauge equation to always be zero but still use the Jacobian of the full system (7.15) and (7.17). If our gauge condition is linear, this procedure would cause us to always remain in the same gauge as the initial guess. However, since (7.17) is nonlinear, we would be committing a change of gauge after each iteration, and therefore we no longer have control over the gauge of the final surface. Nevertheless, we found that if the grid points are evenly distributed across the initial guess surface, then the solver tends to converge to the final surface in a sensible gauge.

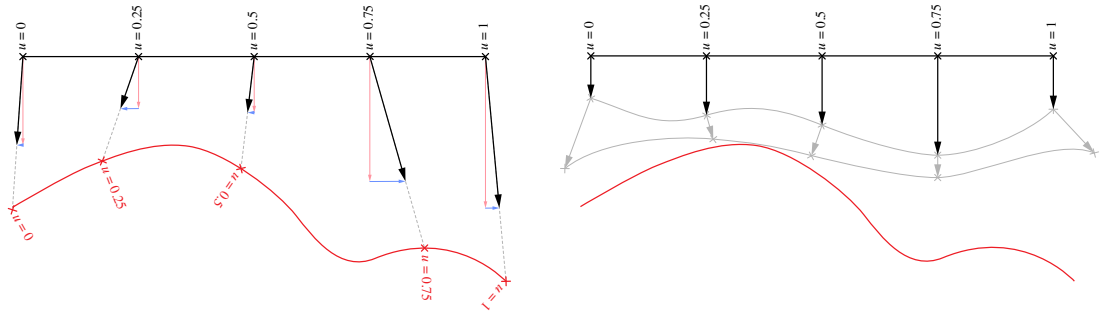

Figure 7.4: Schematic diagrams showing the action of a Newton line search step under different gauge fixing schemes. In both, the horizontal black line represents the initial guess surface, while the red curve represents the solution surface. (left) When a particular $H(u)$ is specified in (7.17), we are not only fixing the red surface as the solution, but also the exact location of the grid points along the solution surface. The direction taken by each line search step (black arrows) consists of both the correction to the actual surface (pink component) and the correction to the gauge condition (blue component). In many cases, this simultaneous gauge correction causes the nonlinear solver to become unstable. (right) By fixing the residual of the gauge equation to always be zero, we no longer specify any target gauge condition on the solution surface. Instead, we are completely projecting out the gauge modes from the Jacobian, leaving the line search direction with no pure gauge component. Since the gauge condition is nonlinear, we are actually committing a slight change of gauge after each Newton step and therefore have no control over the gauge of the final solution.

This 1D case is applicable where the modified cartoon method is used to reduce the simulation to effectively 2+1 dimensions. An example of this is a simulation of the Gregory–Laflamme instability of a 6D Myers–Perry black hole, where the instability is axisymmetric. Following the notation used in our discussion in Section 6.5, we denote coordinates on Σ by (x, y, z, w^A). The additional axisymmetry allows us to also use the modified cartoon method on the (x, y)-plane, thus we can treat y as a cartoon coordinate, leaving only $x^i = (x, z)$ as the grid coordinates on our simulation domain Σ_0. Since $\mathcal{H}$ inherits the same symmetries as Σ, we can locally use w^A

and y as parameters on $\mathcal{H}$ along the reduced dimensions in the neighbourhood of Σ_0. Therefore, $\mathcal{H}$ is parametrised by (u, y, w^A). The coordinate condition (7.17) for u remains unchanged in this setting, and we do not need to provide any additional coordinate conditions for y or w^A. To evaluate (7.15), we use the following relations implied by the symmetries

$$\begin{aligned} \frac{\partial Y}{\partial y} = X \quad , \quad \frac{\partial W^1}{\partial w^2} = Z \quad , \quad \frac{\partial W^2}{\partial w^1} = -Z \quad , \\ \frac{\partial^2 Y}{\partial x \partial y} = \frac{\partial X}{\partial u} \quad , \quad \frac{\partial^2 Z}{\partial w^A \partial w^B} = -\delta_{AB}\, Z \quad , \\ \frac{\partial^2 W^1}{\partial x \partial w^2} = \frac{\partial Z}{\partial u} \quad , \quad \frac{\partial^2 W^2}{\partial x \partial w^1} = -\frac{\partial Z}{\partial u} \quad . \end{aligned} \tag{7.18}$$

We may also generalise this to construct $\mathcal{H}$ with more than one dimension.[1] We begin by recalling the generalised harmonic coordinate condition,

$$\Box_{\mathcal{H}} u^\alpha = \tilde{H}^\alpha(u^\beta). \tag{7.19}$$

Expanding the left-hand side, we get

$$\Box_{\mathcal{H}} u^\alpha = -\gamma^{\beta\gamma}\Gamma^\alpha_{\beta\gamma}. \tag{7.20}$$

To evaluate this, note that u^α can also be used as coordinates on Σ in the neighbourhood of $\mathcal{H}$. We can therefore apply the coordinate transformation rules for the Christoffel symbol

$$\Gamma^\alpha_{\beta\gamma} = \frac{\partial u^\alpha}{\partial x^i}\left(\Gamma^i_{jk}\frac{\partial X^j}{\partial u^\beta}\frac{\partial X^k}{\partial u^\gamma} + \frac{\partial^2 x^i}{\partial u^\beta \partial u^\gamma}\right). \tag{7.21}$$

As it is desirable to avoid having to evaluate the inverse Jacobian $\partial u^\alpha / \partial x^i$ numerically, we consider instead the index-lowered version of this, namely

$$\gamma_{\alpha\delta}\Box_{\mathcal{H}} u^\delta = -\gamma^{\beta\gamma}\gamma_{\alpha\delta}\Gamma^\delta_{\beta\gamma} = \gamma_{ij}\left(\frac{\partial^2 X^i}{\partial u^\beta \partial u^\gamma}\frac{\partial X^j}{\partial u^\alpha} + \Gamma^i_{kl}\frac{\partial X^j}{\partial u^\alpha}\frac{\partial X^k}{\partial u^\beta}\frac{\partial X^l}{\partial u^\gamma}\right). \tag{7.22}$$

[1] The general parametric method for $\mathcal{H}$ with more than one dimension is given here as a theoretical treatment. It has not yet been implemented in code, as the level set approach was sufficient for the published black ring work. The one dimensional case, however, has been successfully deployed in an ongoing project on simulating GL instability in MP black holes, where the new approach is absolutely necessary.

Our general gauge condition is therefore given by

$$\gamma_{ij}\,\gamma^{\beta\gamma}\left(\frac{\partial^2 X^i}{\partial u^\beta \partial u^\gamma}\frac{\partial X^j}{\partial u^\alpha} + \Gamma^i_{kl}\frac{\partial X^j}{\partial u^\alpha}\frac{\partial X^k}{\partial u^\beta}\frac{\partial X^l}{\partial u^\gamma}\right) = H_{(\alpha)}(u^\beta), \tag{7.23}$$

where $H_{(\alpha)}$ are some $d-1$ prescribed functions. Note that in the 1D symmetry-reduced case, this coincides with (7.17), therefore we can regard this as a natural generalisation of our previous result. We can also rewrite the system of (7.15) and (7.23) in a form which makes their common structure manifest:

$$\gamma_{ij}\gamma^{\beta\gamma}\left(\frac{\partial^2 X^i}{\partial u^\beta \partial u^\gamma}S^j + \Gamma^i_{kl}S^j T^k_{(\beta)}T^l_{(\gamma)}\right) = -\left\|S^k\right\|\left(\gamma^{ij} - s^i s^j\right)K_{ij} \tag{7.24}$$

$$\gamma_{ij}\gamma^{\beta\gamma}\left(\frac{\partial^2 X^i}{\partial u^\beta \partial u^\gamma}T^j_{(\alpha)} + \Gamma^i_{kl}T^j_{(\alpha)}T^k_{(\beta)}T^l_{(\gamma)}\right) = H_{(\alpha)}. \tag{7.25}$$

As in the previous 1D case, where the modified cartoon method is used, we can use the w^A coordinates as parameters along the reduced dimensions. In this case, the $\alpha, \beta, \ldots$ indices in the gauge condition (7.25) only range over the non-cartoon parameters, whereas in the physical equation (7.24) they range over all parameters, including the w^A.

7.4 Interpolation of AMR data

In order to find apparent horizons, we must solve a nonlinear PDE on the geometry of each constant time slice Σ. This requires the components of the induced metric γ_{ij} and extrinsic curvature K_{ij} to be C^2 functions on Σ. However, since we have constructed these numerically, their values are only given at the discrete grid points in our AMR hierarchy. In this section, we discuss methods which can be used to interpolate our grid data to obtain sufficiently smooth functions that are suitable for use in further numerical work.

We begin by considering the interpolation problem in one dimension. Let $\{x_0, \ldots, x_N\}$ be our grid points, which we take to be uniformly spaced with separation δx. A scalar function $f(x)$ is sampled onto our grid with values $\{f_0, \ldots, f_N\}$. We wish to construct an interpolant $\hat{f}(x)$ such that $\left|\hat{f}(x) - f(x)\right| = \mathcal{O}(\delta x^n)$ for some chosen n, as $\delta x \to 0$ and $N \to \infty$. Any given point x

can be written as $x = x_i + s\,\delta x$ for some $0 < s < 1$. In order to simplify the implementation, we focus only on interpolation schemes which are linear in the grid data, i.e. those which can be expressed as

$$\hat{f}(x) = \hat{f}(x_i + s\,\delta x) = \sum_{j=-i}^{N-i} f_{i+j}\, w(|s-j|). \tag{7.26}$$

We refer to the function $w(x)$ as the *interpolation kernel*. In order for the interpolant to coincide with the grid data, we require that $w(0) = 1$ and $w(j) = 0$ for each $j \in \mathbb{Z} \setminus \{0\}$. Furthermore, we can obtain an estimate of the dth derivative $f^{(d)}(x)$ via

$$\hat{f}^{(d)}(x) = \hat{f}^{(d)}(x_i + s\,\delta x) = \frac{1}{\delta x^d} \sum_{i=-i}^{n-i} f_i\, w^{(d)}(|s-j|). \tag{7.27}$$

Note that if $\hat{f}(x)$ converges as $\mathcal{O}(\delta x^n)$ then $\hat{f}^{(d)}(x)$ converges as $\mathcal{O}(\delta x^{n-d})$. If $w(x)$ has no support outside of $[-k, k]$, then this can regarded as an application of some $2k$-point stencil to the grid data

$$\hat{f}(x) = \hat{f}(x_i + s\,\delta x) = \sum_{j=-k}^{k} f_{i+j}\, w_j(s). \tag{7.28}$$

One of the most straightforward way to proceed is to use the piecewise *Lagrange polynomials*. In this scheme, we pick the $2k$ grid points nearest to a given point x where an interpolated value is required. We can then fit a unique $(2k-1)$-th degree polynomial on to the corresponding data via

$$\hat{f}(x) = \sum_{j=i-k}^{i+k} f_i \prod_{l \neq j} \frac{x - x_l}{x_j - x_l} \tag{7.29}$$

$$\Rightarrow \quad w_j(s) = \prod_{\substack{l=-k \\ l \neq j}}^{k} \frac{s-l}{j-l}. \tag{7.30}$$

In this scheme, the k-point stencil converges at $\mathcal{O}(\delta x^k)$. It is straightforward to construct an interpolant of arbitrary convergence order in this scheme. Furthermore, [166] provides an efficient pseudocode to generate derivative weights $w_j^{(d)}(s)$ for arbitrarily high d using an arbitrary stencil width. This allows us to leave the stencil width as a fully tunable runtime parameter.

However, the Lagrange polynomial method suffers from a major drawback in that it provides a discontinuous interpolant. The construction above simply fits independent polynomials into each

$(x_i - \delta x/2, x_i + \delta x/2)$ interval without imposing any regularity condition. As a result, $\hat{f}(x)$ will have jump discontinuities halfway between each pair of grid points. To obtain a more suitable interpolant, we can choose to fit an n-th degree polynomial on fewer than $n+1$ data points, and use the remaining degrees of freedom to impose continuity of higher derivatives. This is generically known as *spline interpolation*. Generically, the construction of splines requires the values of derivatives at each grid point, in addition to the function values themselves. An alternative approach is to use the polynomial convolution method [167, 168], where we impose the continuity conditions on the kernel function $w(x)$. This only needs to be done once in advance, and the resulting kernel can then be applied to any dataset. We define a k-point, nth-degree *polynomial convolution kernel* to be the piecewise polynomial

$$p_{k,n}(x) = \begin{cases} a_{1n}\,|x|^n + \ldots + a_{10} & 0 \le |x| < 1 \\ a_{2n}\,|x|^n + \ldots + a_{20} & 1 \le |x| < 2 \\ \vdots & \vdots \\ a_{kn}\,|x|^n + \ldots + a_{k0} & k-1 \le |x| < k \\ 0 & \text{otherwise} \end{cases}, \tag{7.31}$$

where the coefficients a_{ij} are chosen so that the following conditions are satisfied:

- $p_{k,n}(0) = 1$
- $p_{k,n}(i) = 0$ for $i \in \mathbb{Z} \setminus \{0\}$
- $p_{k,n}(x) \in C^{n-2}$

The first two conditions impose k conditions on the cofficients. The third condition imposes a continuity condition for each derivative at each grid point (including the "zeroth derivative"), giving a total of $(n-1)(k+1)$ conditions on the coefficients. However, continuity of even-order derivatives at $x = 0$ is trivial, so this removes $\lfloor n/2 \rfloor$ conditions from the count. This leaves us with a final total of $nk + \lceil n/2 \rceil - 1$ conditions for $k(n+1)$ unknown coefficients. The remaining $k - \lceil n/2 \rceil + 1$ degrees of freedom can be used to match the resulting interpolant $\hat{f}(x_i + s\,\delta x) = \sum_j f_{i+j}\, p_{k,n}(|s-j|)$ to the Taylor expansion of $f(x)$ to as many terms as possible. By solving for all the a_{ij} in this way, the (k,n)-polynomial convolution interpolant can

achieve a convergence rate of $\mathcal{O}(\delta x^{k-\lceil n/2\rceil+3})$. In particular, note that for any given stencil width k, we obtain a smoother interpolant at the cost of a *worse* convergence rate by using higher-degree kernels. Note also that an odd degree kernel $w_{k,2n-1}$ converges at the same rate as the one with next highest even degree $w_{k,2n}$, however the latter will give a smoother result.

Since we need to be able to embed the interpolator within our main GRCHOMBO code, we now focus only on 7-point stencils (i.e., 3 points each side) in order to fit within the ghost region width of our AMR boxes. With the Lagrange polynomial method, we can obtain a $\mathcal{O}(\delta x^6)$ interpolant which is not continuous. The weights are readily determined by expanding out (7.30). For the polynomial convolution method, we have a choice between the piecewise-quartic ($n = 4$) kernel which is C^2 and converges at $\mathcal{O}(\delta x^4)$, or a piecewise-quintic ($n = 5$) kernel which is C^3 but converges at a slower rate $\mathcal{O}(\delta x^3)$. Since the apparent horizon problem only requires first derivatives, the quartic kernel should be sufficiently smooth and gives a faster convergence rate. For completeness, we give the stencil weights for both here.

	Quartic Convolution	Quintic Convolution
$w_{-2}(s)$	$\frac{1}{12}\, s\,(1-s)^3$	$\frac{3}{64}\, s\,(1-s)^4$
$w_{-1}(s)$	$-\frac{1}{12}\, s\,(1-s)\,(2-s)\,(4-3s)$	$-\frac{1}{64}\, s\,(1-s)\left(38-42s-6s^2+13s^3\right)$
$w_0(s)$	$\frac{1}{6}\,(1-s)\left(6+6s-9s^2+s^3\right)$	$\frac{1}{32}\,(1-s)\left(32+32s-36s^2-36s^3+27s^4\right)$
$w_1(s)$	$\frac{1}{6}\, s\left(4+9s-6s^2-s^3\right)$	$\frac{1}{32}\, s\left(19+40s+18s^2-72s^3+27s^4\right)$
$w_2(s)$	$-\frac{1}{12}\, s\,(1-s)\,(1+s)\,(1+3s)$	$-\frac{1}{64}\, s\,(1-s)\left(3+15s+33s^2-13s^3\right)$
$w_3(s)$	$\frac{1}{12}\, s^3\,(1-s)$	$\frac{3}{64}\, s^4\,(1-s)$

In practice, since the apparent horizon usually lies within the finest levels in the AMR hierarchy, the grid data to be interpolated is available at extremely high resolutions. The Lagrange polynomial and the two convolution kernels given above yield very similar end results. However, since we employ Newton's method with numerically approximated Jacobian, the lack of continuity in the Lagrange interpolant tends to lead to a less robust code. As only first derivatives are required in our computation here, the quartic convolution method provides a good compromise between convergence rate and smoothness for our purpose.

Chapter 8

End Point of 5D Black Ring Instabilities

This chapter provides a brief summary of the results of a collaborative work, as published in [169]. Further work on this topic is currently ongoing by other members of the collaboration. This chapter is included for completeness, as an example of a successful application of the tools and techniques discussed in the previous two chapters. The analysis methods described in Section 8.3 are entirely my own work.

8.1 Introduction

Black rings are asymptotically flat and stationary black hole solutions to the D-dimensional vacuum Einstein equation whose Killing horizon has an $S^{D-3} \times S^1$ topology. This family of solutions was first found in $D = 5$, where the metric is known analytically [14]. Five dimensional black rings are parametrised by two parameters: the radius $R > 0$ and the thickness $0 < \nu < 1$. Solutions with $\nu < 1/2$ are called *thin*, while those with $\nu > 1/2$ are called *fat*, where these labels reflect the geometry of the ring's horizon in the obvious manner. As previously discussed in Section 6.1, the resemblance between thin black rings and boosted Kaluza-Klein black strings led to the speculation that thin rings could also be subject to the Gregory–Laflamme (GL) instability [135]. The authors of [135] also constructed a model which approximates thin rings as boosted black strings to show that this should indeed be the case, however their argument is only valid when the ring is very thin. More recently, linearised perturbation analysis of black rings [136] found axisymmetry-breaking unstable modes which affect not only *all* thin rings, but also some slightly fat ones as well. These modes deform the S^2 so as to cause the ring's thickness to become non-uniform in a qualitatively

similar manner to the GL modes in black strings.

The numerical simulation of black strings presented in [31] showed that the GL instability results in a self-similar process, by which the string condenses into a sequence of spherical droplets adjoined by ever thinner black strings. The geometric decay in time elapsed between successive generations of this process led the authors to conclude that the adjoining strings would shrink to zero radius in finite time, causing the horizon to pinch off to form disconnected pieces. This singular topology-changing process is regarded as evidence of violation of the weak cosmic censorship conjecture in a 5D, asymptotically Kaluza-Klein setting. Demonstrating that the same process occurs in thin black rings could therefore provide an example of a violation of the weak cosmic censorship conjecture in an asymptotically flat spacetime, which is more closely related to the astrophysical setting. Resolution of the weak cosmic censorship conjecture is one of the most important open problems in general relativity, as it must hold true if GR is to be a complete, predictive classical theory.

Aside from the axisymmetry-breaking GL instability, we could also ask how black rings react to radial perturbations. To study this, [122] considered constraint-violating deformations to the black ring geometry. This 'off-shell deformation' can be regarded as an application of an external radial force to the black ring. Fat rings ($\nu > 1/2$) appear at a local maximum of the effective potential of the force, while the thin rings ($\nu < 1/2$) appear at a local minimum. The authors argued that local maxima correspond to a configuration where the inner S^1 of the ring requires an inward-pulling tension to restore equilibrium, suggesting that the ring would tend to collapse in this setting, with the hole in the middle closing to form a Myers-Perry black hole. The existence of a radial instability in fat rings is confirmed in [170] through the application of local Penrose inequalities.

To gain an understanding of the nonlinear dynamics and the end states of these various instabilities, we conduct a numerical simulation of perturbed black rings using the tools and techniques discussed in Chapters 6 and 7. Our simulations reveal a new family of non-axisymmetric instability which had not been identified prior to our results. Unlike the GL instability, this does not cause a substantial variation in thickness along the S^1. Instead, it acts to deform the shape of the S^1 itself, stretching certain regions apart while bringing other regions closer to the centre. This new instability affects both thin and fat rings. In keeping with the terminology used in blackfold analyses [28], we call it the *elastic* instability. Unless the ring is extremely thin, this elastic mode causes the ring to merge

in the middle, forming a topologically spherical horizon before the GL mode has time develop its expected singular behaviour.

In this chapter, we will give a brief summary of our simulation methods and results, in line with those reported in [169]. The project to numerically investigate the dynamics of black ring instabilities is an ongoing collaborative work, and further detailed results will be reported by other members of the collaboration in future publications.

8.2 Numerical setup

The metric of a stationary black ring with radius R and thickness ν can be expressed in terms of ring coordinates via

$$\begin{aligned} \mathrm{d}s^2 = &-\frac{F(y)}{F(x)}\left(\mathrm{d}t - C\,R\,\frac{1+y}{F(y)}\,\mathrm{d}\psi\right)^2 \\ &+\frac{R^2\,F(x)}{(x-y)^2}\left[\frac{\mathrm{d}x^2}{G(x)} - \frac{\mathrm{d}y^2}{G(y)} + \frac{1}{1+\nu^2}\left(\frac{G(x)}{F(x)}\,\mathrm{d}\phi^2 - \frac{G(y)}{F(y)}\,\mathrm{d}\psi^2\right)\right] \end{aligned} \tag{8.1}$$

$$F(\xi) := 1 + \frac{2\,\nu\,\xi}{1+\nu^2} \quad , \quad G(\xi) := \left(1-\xi^2\right)(1+\nu\,\xi) \quad , \quad C := \sqrt{\frac{2\,\nu\,(1+\nu)^3}{(1-\nu)\left(1+\nu^2\right)^3}}.$$

Note that the above line element is given in the form which manifestly has no conical singularities, and whose angular coordinates have a canonical period, $\Delta\phi = \Delta\psi = 2\pi$. The range of the x and y coordinates are $-1 \leq x \leq 1$ and $-1/\nu < y < -1$. The metric clearly has a coordinate singularity at the Killing horizon $y = -1/\nu$. To remove this singularity, we recast (8.1) into isotropic coordinates using the transformation introduced in [170],

$$y(z) := z^2 - 1/\nu\ . \tag{8.2}$$

We then simply take the spacelike hypersurface Σ to be the constant-t slice. The induced metric on Σ is then given by

$$\begin{aligned} \mathrm{d}s^2_\Sigma = &-\frac{C^2\,R^2\,(1+y(z))^2}{F(x)\,F(y(z))}\,\mathrm{d}\psi^2 \\ &+\frac{R^2\,F(x)}{(x-y(z))^2}\left[\frac{\mathrm{d}x^2}{G(x)}-\frac{4\,\mathrm{d}z^2}{1-y(z)^2}+\frac{1}{1+\nu^2}\left(\frac{G(x)}{F(x)}\,\mathrm{d}\phi^2-\frac{G(y(z))}{F(y(z))}\,\mathrm{d}\psi^2\right)\right]. \end{aligned} \tag{8.3}$$

We apply a perturbation to the analytic ring initial data by modifying only the conformal factor χ. A generic perturbation can be decomposed into Fourier modes as

$$\chi = \chi_0 \sum_{m=0}^{\infty} f_m(r_2)\,\cos m\psi, \tag{8.4}$$

where $r_2 := \sqrt{X^2+Y^2}$ and $\psi = \tan^{-1} Y/X$. For the work presented in this chapter, we focus entirely on $m = 2$ perturbations. In terms of Cartesian coordinates, these take the general form

$$\chi = \chi_0\left[1+\varepsilon\,F\big(X^2+Y^2, X^2-Y^2, Z\big)\right], \tag{8.5}$$

where ε is the perturbation amplitude. The specific perturbation function used in our work is given by

$$F = \frac{1}{\big(1+\Xi^2\big)^{3/2}}\,\frac{X^2-Y^2}{\Sigma}$$
$$\Sigma := \sqrt{\big(\tilde{R}^2+Z^2+r_2^2\big)^2-4\,\tilde{R}\,r_2^2} \quad , \quad \Xi := \frac{4\,\Sigma\,(1-\nu)}{\nu\,\big(r^2+\tilde{R}^2-\Sigma\big)}.$$

This choice of perturbation violates both the Hamiltonian and momentum constraints. However, for small ε the constraint-damping property of CCZ4 can quickly suppress any initial constraint violations. Indeed, we can consider this initial phase of the CCZ4 evolution as a relaxation scheme for solving the initial constraint equations.

8.3 Analysing the black ring's horizon geometry

In order to measure the extent of the GL instability, we need to characterise the size of the transverse S^2 spheres along the ring's horizon. For the stationary axisymmetric solution, the spheres are obtained simply by fixing the angular coordinate ψ constant. However, for our dynamical horizon, taking a constant-ψ section is no longer a sensible choice, as the axisymmetry is broken. Generally speaking, there is no preferred way to slice a transverse section through an arbitrarily chosen point on $\mathcal{H}$. One may contemplate using spacelike geodesics along $\mathcal{H}$ to define a preferred S^2, however the section obtained by slicing along such geodesics will generically not form a smooth S^2. Our approach here is to instead extremise the sphere's area over all the possible smooth S^2 that can be sliced on $\mathcal{H}$. Specifically, consider an apparent horizon $\mathcal{H}$ with topology $S^2 \times S^1$, which is defined as the level set $\rho = F(\xi, \psi)$ as previously described. The line element on $\mathcal{H}$ is given by

$$\mathrm{d}s^2_{\mathcal{H}} = h_{\xi\xi}\,\mathrm{d}\xi^2 + h_{\psi\psi}\,\mathrm{d}\psi^2 + 2\,h_{\xi\psi}\,\mathrm{d}\xi\,\mathrm{d}\psi + Z^2\,h_{ww}\,\mathrm{d}\phi^2 \tag{8.6}$$

$$Z := \frac{R\,\sin\xi}{\cos\xi + 1/\tanh\rho}, \tag{8.7}$$

where $h_{..}$ denote the components of the metric on Σ pulled back onto $\mathcal{H}$. Recall also that we allow $R = R(\psi)$ in order to accommodate stretched horizons.

We parametrise an S^2 section on $\mathcal{H}$ via $\psi = \Psi(\xi)$. The line element on the S^2 is given by

$$\mathrm{d}s^2_{S^2} = \left(h_{\psi\psi}\,\Psi'(\xi)^2 + 2\,h_{\xi\psi}\,\Psi'(\xi) + h_{\xi\xi}\right)\mathrm{d}\xi^2 + Z^2\,h_{ww}\,\mathrm{d}\phi^2, \tag{8.8}$$

and thus its area is given by

$$A_{S^2} = 4\pi \int_{-\pi}^{\pi} \mathrm{d}\xi \sqrt{\left(h_{\psi\psi}\,\Psi'(\xi)^2 + 2\,h_{\xi\psi}\,\Psi'(\xi) + h_{\xi\xi}\right) Z^2\,h_{ww}}. \tag{8.9}$$

To extremise the area, we vary A_{S^2} with respect to $\Psi(\xi)$ to obtain the Euler-Lagrange equation, which schematically takes the form

$$2\,Z\left(h_{\xi\xi}\,h_{\psi\psi} - h^2_{\xi\psi}\right)\Psi''(\xi) + A\,\Psi'(\xi)^3 + B\,\Psi'(\xi)^2 + C\,\Psi'(\xi) + D = 0, \tag{8.10}$$

where A, B, C, and D are some complicated expressions involving only the metric components, Z,

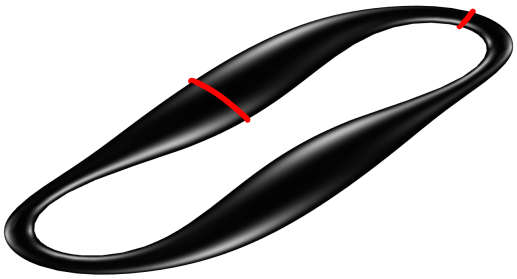

Figure 8.1: The red curve shows the transverse S^2 with minimum and maximum area of a particularly deformed ring horizon, as found by the variational method.

and their derivatives (i.e., independent of $\Psi(\xi)$). In practice, the exact expression can be produced using computer algebra. The boundary conditions are obtained by requiring that the solutions are smooth spheres, in particular that there are no singularities at the poles. Expanding (8.8) near $\xi = 0$, we can see that the smoothness condition is given by

$$h_{\psi\psi}\,\Psi'(0)^2 + 2\,h_{\xi\psi}\,\Psi'(\xi) + h_{\xi\xi} = \frac{R^2 h_{ww}}{1 + 1/\tanh\rho}, \tag{8.11}$$

which can be solved to obtain a Neumann boundary condition

$$\Psi'(0) = \frac{1}{h_{\psi\psi}}\left(-h_{\xi\psi} + \sqrt{h_{\xi\psi}^2 - h_{\psi\psi}\left(h_{\xi\xi} - \frac{R^2\,h_{ww}}{1 + 1/\tanh\rho}\right)}\right). \tag{8.12}$$

Note that the choice of positive square root is dictated by the direction of the black ring's rotation. Likewise, by considering (8.8) near $\xi = \pi$, we get

$$\Psi'(\pi) = \frac{1}{h_{\psi\psi}}\left(-h_{\xi\psi} + \sqrt{h_{\xi\psi}^2 - h_{\psi\psi}\left(h_{\xi\xi} - \frac{R^2\,h_{ww}}{-1 + 1/\tanh\rho}\right)}\right). \tag{8.13}$$

The resulting boundary value problem can be discretised using finite differences and solved via Newton's method. In general, we expect there to be multiple S^2 solutions corresponding to local extrema and saddle points. For example, with the $m = 2$ perturbation, the symmetry of the system is such that we should expect the minimal and maximal S^2 to have another copy on the opposite side of the ring, giving a total of four solutions during the first generation of GL instability. Since solutions obtained by Newton's method depend on the initial guess used, we proceed by running the Newton solver on multiple initial guesses around the ring in order to converge to all possible

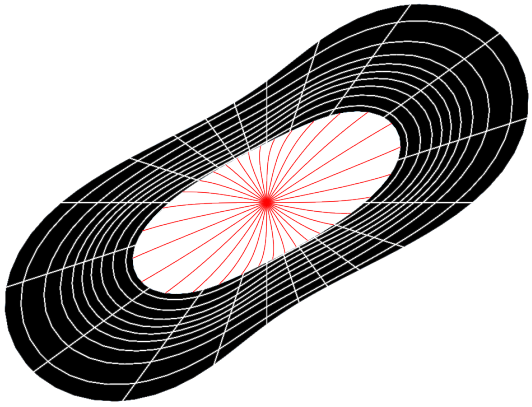

Figure 8.2: To measure elastic deformation of the ring horizon, we use distances along spacelike geodesics from the centre of the ring to the point where they intersect the inner S^1 of the horizon. These are shown by the red curves.

solutions. A good choice for this is to use constant-ψ curves as initial guesses. Even though these curves do not satisfy our boundary conditions, we found that the Newton solver can quickly correct for this. We can then pick out the solutions with minimal and maximal areas, as displayed in Figure 8.1, to characterise the GL instability.

Over the course of our simulation of thin black rings, we identified a new type of non-axisymmetric instability which deforms the shape of the ring but does not cause its thickness to vary like in GL. This *elastic mode* instability is somewhat more mathematically straightforward to characterise compared to the GL instability. Let (r, ψ) be the polar coordinates on the $z = 0$ plane. For each $0 \leq \psi_0 < 2\pi$, let $(r(\lambda), \psi(\lambda))$ be the affinely-parametrised spacelike geodesic emanating from the origin in the direction ψ_0, i.e. $r(0) = 0$, $\psi(0) = \psi_0$ and $\psi'(0) = 0$. The geodesic intersects $\mathcal{H}$ at some parameter $\lambda = \lambda_*$ with $\psi(\lambda_*) = \psi_*$, and we denote length from $\lambda = 0$ to $\lambda = \lambda_*$ by $\ell(\psi_0)$. By treating ψ_* as a bijective function of ψ_0, we can use $\ell(\psi_*^{-1}(\psi_*))$ to characterise the distance from the centre of the ring to various points along its inner S^1. Any variation in this function is indicative of the deformation in the shape of the ring, which is precisely what we are after. This procedure is depicted in Figure 8.2.

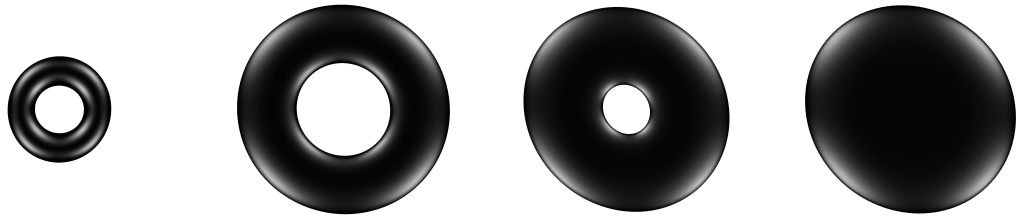

$t/\sqrt{M} = 0$ $t/\sqrt{M} = 5.53$ $t/\sqrt{M} = 17.3$ $t/\sqrt{M} = 17.4$

Figure 8.3: Evolution of a fat ring with $\nu = 0.7$, subject to an $m = 2$ perturbation mode with amplitude $A = 0.02$. The ring initially grows in coordinate size as the evolution goes through the gauge adjustment phase. This ceases at $t/\sqrt{M} \sim 5$. We can then clearly see the radial instability kicking in, causing the ring to become thicker over time. Eventually, the hole in the middle closes at $t/\sqrt{M} \sim 17$ and the geometry settles down to that of a Myers–Perry black hole. Here, we show snapshots corresponding to the coarsest time steps immediately before and after a spherical horizon is formed. Note that there is no substantial deformation in the shape of the ring's S^1 or non-uniformity in the S^2 size in this case.

8.4 Summary of results

For very fat rings, the evolution is dominated by a radial instability, as predicted by [122, 170]. When such a ring is subject to a non-axisymmetric perturbation, the size of its S^2 increases in a roughly uniform manner, while the shape of the S^1 does not deform significantly. This eventually leads to the closure of the hole in the middle, and the final state here is therefore a Myers–Perry black hole. Snapshots of the apparent horizon showing the evolution of a $\nu = 0.7$ fat ring subject to an $m = 2$ perturbation is shown in Figure 8.3.

For rings whose thickness lies in the range $0.4 \lesssim \nu \lesssim 0.6$, the dynamics of the instability resulting from an initial $m = 2$ perturbation are dominated by a deformation of the S^1. As shown in the second snapshot of Figure 8.4, this mode of instability stretches the ring without causing a substantial variation in thickness. Prior to the publication of our work, this instability of the black ring had not previously been identified in the literature as far as we are aware. We call this the *elastic mode instability*, in keeping with the terminology used in blackfolds work [28]. This deformation eventually results in the ring merging in the middle along the short side of the stretch, forming a topologically spherical horizon with an elongated shape. This elongated black hole then proceeds to shed its angular momentum before settling down to the round Myers–Perry black hole

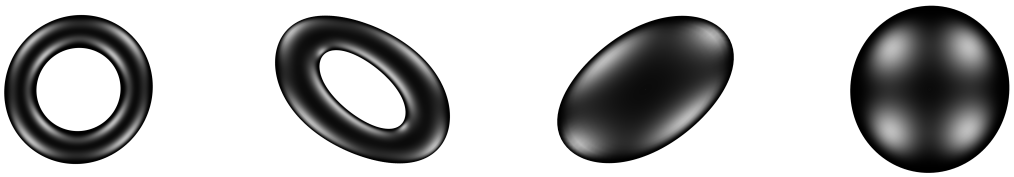

$t/\sqrt{M} = 2.93$ $t/\sqrt{M} = 10.4$ $t/\sqrt{M} = 12.2$ $t/\sqrt{M} = 32.6$

Figure 8.4: Evolution of a ring at the cusp between the fat and thin branch with $\nu = 0.5$, subject to an $m = 2$ perturbation mode with amplitude $A = 0.02$. The first snapshot shows the ring's horizon after gauge adjustment. The evolution is clearly dominated by a stretching behaviour of the S^1, while the S^2 size still remains roughly uniform throughout. We call this an *elastic mode instability*. This instability eventually causes the ring to merge along the short side of the stretch, forming an elongated black hole with a topologically spherical horizon. The black hole then quickly sheds angular momentum before settling down to a round, Myers–Perry final state.

as a final state.

For thinner rings with $0.2 \lesssim \nu \lesssim 0.4$, the GL behaviour starts to become important, and the resulting dynamics is a nonlinear combination of the GL and elastic modes. In Figures 8.5-8.6, we can see that the ring's thickness becomes markedly non-uniform in the advanced stages of the evolution, while the S^1 also becomes highly stretched. In this regime, we observed that the stretching behaviour causes a topologically spherical horizon to form before the GL mode can develop to the point where the horizon pinches off. Therefore, we do not observe any potential violation of weak cosmic censorship here.

We can also clearly see the presence of both the elastic and GL instabilities through geometric measurements, as described in Section 8.3. In Figure 8.7(a), the plot shows a divergence between the size of the largest and the smallest S^2 of the ring. This significant non-uniformity in the ring's thickness is consistent with the expected behaviour of GL dynamics. On the other hand, 8.7(b) also shows a divergence in the shortest and longest S^1 radii of the ring, which is precisely the behaviour of an elastic mode instability. We can also observe the S^1 radii eventually plunging towards zero, showing the hole in the middle of the ring closing before further generations of GL bulges can develop.

For very thin rings with $\nu \lesssim 0.15$, the GL instability becomes dominant in the nonlinear dynamics,

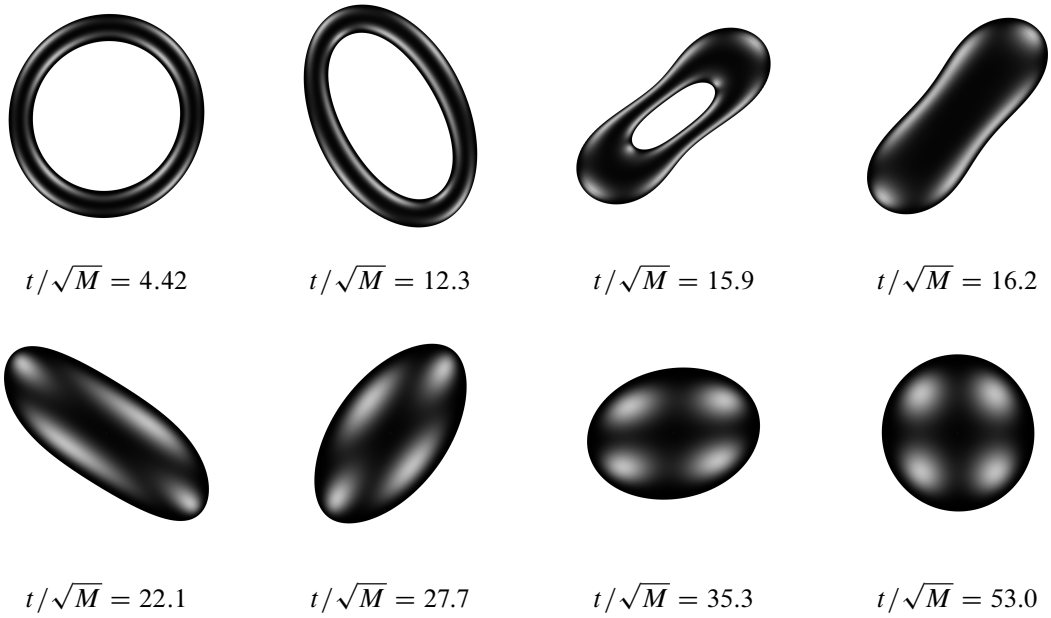

Figure 8.5: Evolution of a thin ring with $\nu = 0.3$, subject to an $m = 2$ perturbation mode with amplitude $A = 0.002$. The first snapshot shows the ring's horizon after gauge adjustment. A GL instability is clearly present here, as the S^2 becomes pronouncedly non-uniform. However, the $m = 2$ elastic mode remains important and the ring still merges along the short side of the stretch. The final state is therefore still a Myers–Perry black hole.

Figure 8.6: Evolution of a thin ring with $\nu = 0.2$, subject to an $m = 2$ perturbation mode with amplitude $A = 0.0005$. The first snapshot shows the ring's horizon after gauge adjustment. The effect of the GL instability is clearly more substantial here compared to the $\nu = 0.3$ case, while the elastic stretching behaviour remains important. Unlike in the previous case, the fatter part of the ring occurs along the short side of the stretching. The apparent horizon finder fails after the final snapshot shown above, and we suspect that this is because the horizon first goes through a brief intermediate stage with the topology of a two-holed torus. Finding an appropriate parametrisation of such a genus-2 horizon remains a challenging problem.

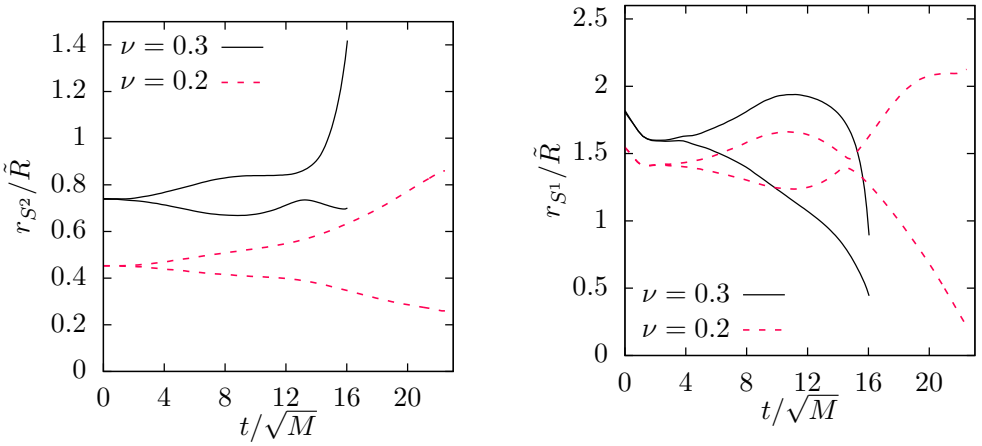

(a) Evolution of the sizes of the largest and smallest S^2 of the ring.

(b) Evolution of the shortest and longest geodesic distances from the middle of the ring to its inner S^1.

Figure 8.7: Geometric measurements of the apparent horizons for rings with $\nu = 0.2$ and $\nu = 0.3$. The methods by which these quantities are calculated are described in Section 8.3. The divergence between the largest and smallest S^2 size is the result of a GL instability. The initial divergence between the largest and smallest S^1 radii is the result of an elastic mode behaviour. The S^1 radii eventually plunge towards zero as the hole in the middle closes, and a spherical horizon forms.

as shown in Figure 8.8. Although we only explicitly perturbed the ring with an $m = 2$ mode, various numerical errors and noises also act to excite all angular modes throughout the evolution. The amplitudes of the perturbations induced by these noises are orders of magnitude lower than our explicit perturbation, however they could nevertheless become important if modes other than $m = 2$ have a faster-growing instability. This appears to be the case at $\nu = 0.15$, where the evolution clearly exhibits an $m = 4$ behaviour. The $m = 4$ elastic mode stretches the ring's S^1 along two perpendicular directions, resulting in a square-shaped deformation without a particular 'short side'. In Figure 8.8, the effect of the initial $m = 2$ perturbation is still visible, as two of the GL bulges are large than the others. As the evolution continues, the ring begins to stretch in such a way as to cause difficulties for our apparent horizon finder. However, we can use a contour of the conformal factor χ as a proxy to determine the qualitative shape of the horizon. As shown in Figure 8.9, a second generation of bulges appear to have developed along the thinner necks of the ring's horizon. This is the same behaviour as observed in the black string in [31], and we expect that our evolution will continue to develop smaller bulges connected by ever thinner necks over time. As we did not fine-tune our initial condition, this behaviour can be regarded as generic, and therefore provides evidence that weak cosmic censorship is violated in the neighbourhood of very

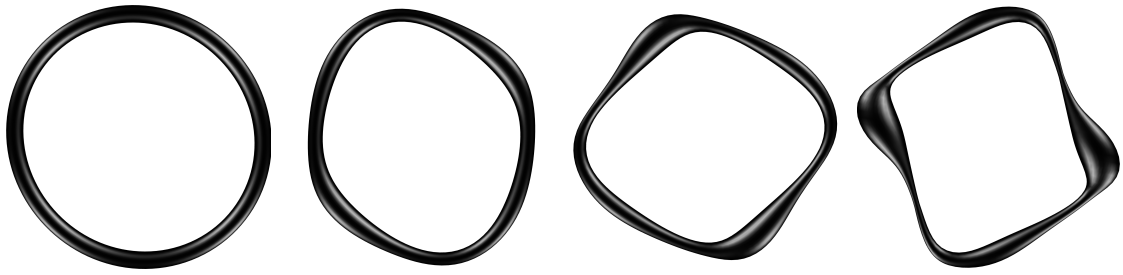

Figure 8.8: Evolution of a very thin ring with $\nu = 0.15$. The first snapshot shows the ring's horizon after gauge adjustment. Although we only explicitly perturbed the ring with an $m = 2$ mode with amplitude $A = 0.00005$, the dynamics becomes dominated by an $m = 4$ instability excited by various numerical noises. The GL instability is very prominent here, and the ring's horizon has a clear structure of bulges connected by thin necks. The $m = 4$ elastic mode deforms the ring in a counter-balanced fashion without a preferred "short side" as in previous cases, allowing more time for the GL instability to fully develop before the ring collapses. This type of behaviour is likely to cause the ring horizon to pinch off in finite time, leading to a violation of weak cosmic censorship.

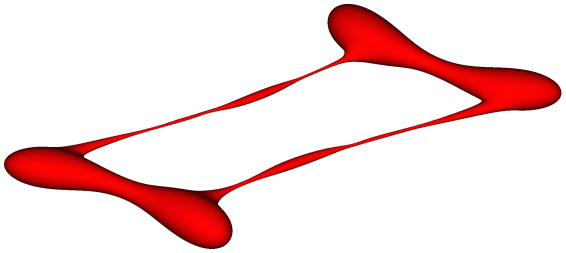

Figure 8.9: The contour $\chi = 0.2$ of the same $\nu = 0.15$ ring as shown in Figure 8.8. The unusual geometry of the ring poses a difficulty for the apparent horizon finder, but the contour displayed above provides a qualitative shape of the horizon.

thin rings.

In Figure 8.10, we show the result of an evolution of the same $\nu = 0.15$ ring without any added initial perturbation. Here, we obtain a 'cleaner' evolution, where the entire dynamics is governed by the $m = 4$ modes excited by the grid noises, resulting in four virtually identical GL bulges connected by extremely thin necks. Without the additional stretching effect from the $m = 2$ perturbation, however, the time scale required for the necks to become sufficiently thin for us to observe further GL behaviour is much longer than in the previous case. Therefore, significant additional computational resources will be required to reach the second GL generation in this case.

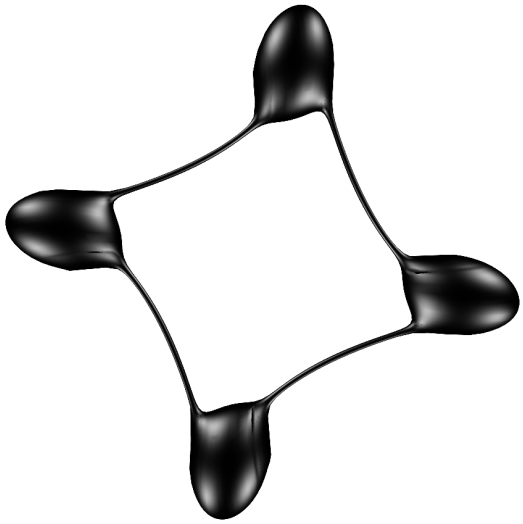

Figure 8.10: Apparent horizon at $t/\sqrt{M} = 32.1$ of a very thin ring with $\nu = 0.15$, without any explicitly added perturbation. The horizon clearly exhibits behaviour which shows that the dynamics is dominated by an $m = 4$ mode instability. All four first-generation bulges have the same size and shape, and the necks become extremely thin compared to the bulges.

Finally, we can compare the growth rate and frequency of the gravitational wave emitted from the black rings in our simulation with the perturbative results of [136]. We do this by measuring the quantity

$$h_+ := \frac{\tilde{\gamma}_{xx} - \tilde{\gamma}_{yy}}{2} \left(\frac{r}{\tilde{R}}\right)^{3/2} \tag{8.14}$$

along the z-axis, far away from the ring's horizon. Since we expect the dynamics of thin rings to contain GL-mode and elastic-mode behaviours, we fit the raw wave form to the model

$$A_1 \sin(\Re\varpi_1 t + \varphi_1)\, e^{\Im\varpi_1 t} + A_2 \sin(\Re\varpi_2 t + \varphi_2)\, e^{\Im\varpi_2 t}, \tag{8.15}$$

where $\{A_1, A_2, \varphi_1, \varphi_2, \varpi_1, \varpi_2\}$ are chosen to minimise the difference between the model and the data. The results are displayed in Figure 8.11, where our data for the GL mode are in excellent agreement with the results of [136] over the range $0.15 \leq \nu \leq 0.4$.

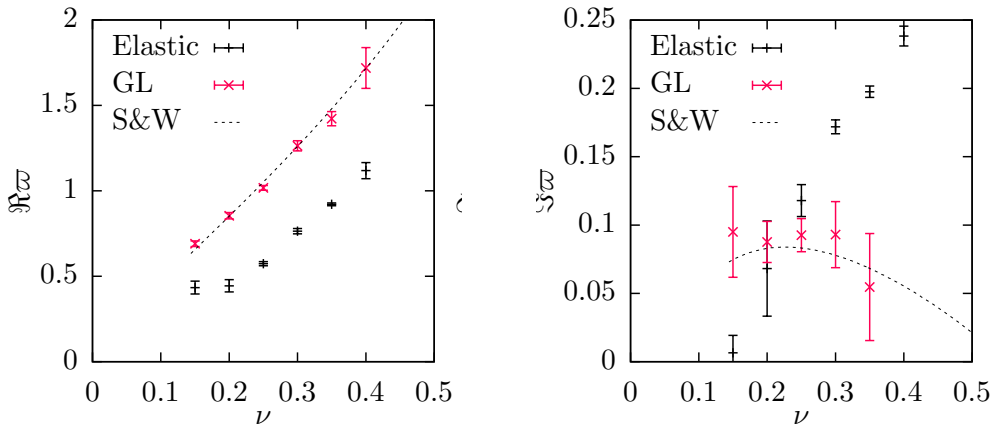

(a) Frequency

(b) Growth rate

Figure 8.11: Growth rate and frequency of the gravitational wave emitted from thin rings subjected to an $m = 2$ perturbation. The perturbative results of [136] are shown as dotted curves.

CHAPTER 9

Summary

In this dissertation, we explored the application of numerical methods to construct various solutions to the Einstein equation which lie outside the usual astrophysical setting. The work presented can be divided into two distinct categories. In the first part of the dissertation—chapters 1-5—we considered the construction of stationary solutions which are physically motivated through the AdS/CFT correspondence. This has allowed us to investigate thermodynamic aspects of certain strongly coupled conformal field theories in various physical configurations. The 'rotating black droplets' of chapter 3 allowed us to study Hawking radiation in the background of a rotating black hole. The 'plasma balls' of chapter 4 allowed us to study a metastable, localised ball of deconfined CFT plasma surrounded by a confining vacuum. Finally, the construction of AdS black rings in chapter 6 provided us with a more complete picture of the phase diagram of AdS black holes in 5D, which also correspond to the thermodynamic phase diagram of $\mathcal{N} = 4$ SYM in 4D. The programme to numerically construct new stationary spacetimes with a negative cosmological constant remains a very active area of research.

The second part of the dissertation—chapters 6-8—concerned the dynamical evolution of higher dimensional spacetimes. This line of work requires the use of dimensional reduction and adaptive mesh refinement techniques, which were implemented in our GRCHOMBO numerical relativity code. The extended singularities which arise in higher dimensional black holes led us to introduce a new artificial viscosity term in order to regulate them over the course of the evolution. The unusual and complex geometries arising from black hole instabilities in higher dimensions also required new techniques to be developed in order to locate the apparent horizon over the course of the evolution, which were discussed in chapter 7. These tools and techniques were then applied

to evolve perturbed asymptotically flat black rings in five dimensions. In addition to the expected Gregory–Laflamme instability, our simulations also revealed a new type of unstable mode which acts to deform the shape of the ring's S^1. For very thin rings, the GL instability eventually unfolds in a manner which is consistent with the behaviour previously observed in Kaluza–Klein black strings, namely the development of a family of bulges connected by ever thinner string-like 'necks' which eventually reduce to zero size. We interpret this as the first concrete evidence of a possible violation of the weak cosmic censorship conjecture in the asymptotically flat setting. Further work is currently being carried out by other members of the collaboration to apply these tools and techniques to investigate instabilities of ultraspinning Myers–Perry black holes in six dimensions, where it is expected that both the GL instability and 'bar-mode' instability would also lead to violations of weak cosmic censorship.

References

[1] M. W. Choptuik, "Universality and scaling in gravitational collapse of a massless scalar field," *Phys. Rev. Lett.* **70** (1993) 9–12.

[2] M. Choptuik, "The binary black hole grand challenge project," in *Computational Astrophysics; 12th Kingston Meeting on Theoretical Astrophysics*, vol. 123, p. 305. 1997.

[3] F. Pretorius, "Evolution of binary black hole spacetimes," *Phys. Rev. Lett.* **95** (2005) 121101, arXiv:gr-qc/0507014.

[4] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, "Accurate evolutions of orbiting black-hole binaries without excision," *Phys. Rev. Lett.* **96** (2006) 111101, arXiv:gr-qc/0511048.

[5] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter, "Gravitational wave extraction from an inspiraling configuration of merging black holes," *Phys. Rev. Lett.* **96** (2006) 111102, arXiv:gr-qc/0511103.

[6] **Virgo, LIGO Scientific** Collaboration, B. P. Abbott *et al.*, "Observation of Gravitational Waves from a Binary Black Hole Merger," *Phys. Rev. Lett.* **116** no. 6, (2016) 061102, arXiv:1602.03837.

[7] **Virgo, LIGO Scientific** Collaboration, B. P. Abbott *et al.*, "GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence," *Phys. Rev. Lett.* **116** no. 24, (2016) 241103, arXiv:1606.04855.

[8] W. Israel, "Event horizons in static vacuum space-times," *Phys. Rev.* **164** (1967) 1776–1779.

[9] B. Carter, "Axisymmetric Black Hole Has Only Two Degrees of Freedom," *Phys. Rev. Lett.* **26** (1971) 331–333.

[10] S. W. Hawking, "Black holes in general relativity," *Commun. Math. Phys.* **25** (1972) 152–166.

[11] D. C. Robinson, "Uniqueness of the Kerr black hole," *Phys. Rev. Lett.* **34** (1975) 905–906.

[12] R. P. Kerr, "Gravitational field of a spinning mass as an example of algebraically special metrics," *Phys. Rev. Lett.* **11** (1963) 237–238.

[13] R. C. Myers and M. Perry, "Black Holes in Higher Dimensional Space-Times," *Annals Phys.* **172** (1986) 304.

[14] R. Emparan and H. S. Reall, "A Rotating black ring solution in five-dimensions," *Phys.Rev.Lett.* **88** (2002) 101101, arXiv:hep-th/0110260.

[15] A. Pomeransky and R. Sen'kov, "Black ring with two angular momenta," arXiv:hep-th/0612005.

[16] H. Elvang and P. Figueras, "Black Saturn," *JHEP* **0705** (2007) 050, arXiv:hep-th/0701035.

[17] H. Iguchi and T. Mishima, "Black di-ring and infinite nonuniqueness," *Phys.Rev.* **D75** (2007) 064018, arXiv:hep-th/0701043.

[18] J. Evslin and C. Krishnan, "The Black Di-Ring: An Inverse Scattering Construction," *Class.Quant.Grav.* **26** (2009) 125018, arXiv:0706.1231.

[19] K. Izumi, "Orthogonal black di-ring solution," *Prog.Theor.Phys.* **119** (2008) 757–774, arXiv:0712.0902.

[20] H. Elvang and M. J. Rodriguez, "Bicycling Black Rings," *JHEP* **0804** (2008) 045, arXiv:0712.2425.

[21] B. Kleihaus, J. Kunz, and E. Radu, "Black rings in six dimensions," *Phys. Lett.* **B718** (2013) 1073–1077, arXiv:1205.5437.

[22] B. Kleihaus, J. Kunz, and E. Radu, "Black Rings in More Than Five Dimensions," in *Proceedings, 13th Marcel Grossmann Meeting on Recent Developments in Theoretical and Experimental General Relativity, Astrophysics, and Relativistic Field Theories (MG13): Stockholm, Sweden, July 1-7, 2012*, pp. 1321–1323. 2015.

[23] G. J. Galloway and R. Schoen, "A Generalization of Hawking's black hole topology theorem to higher dimensions," *Commun. Math. Phys.* **266** (2006) 571–576, arXiv:gr-qc/0509107.

[24] J. M. Maldacena, "The Large N limit of superconformal field theories and supergravity," *Adv.Theor.Math.Phys.* **2** (1998) 231–252, arXiv:hep-th/9711200.

[25] V. A. Belinsky and V. E. Zakharov, "Integration of the Einstein Equations by the Inverse Scattering Problem Technique and the Calculation of the Exact Soliton Solutions," *Sov. Phys. JETP* **48** (1978) 985–994. [Zh. Eksp. Teor. Fiz.75,1953(1978)].

[26] V. Belinsky and V. Sakharov, "Stationary Gravitational Solitons with Axial Symmetry," *Sov.Phys.JETP* **50** (1979) 1–9.

[27] H. S. Reall, "Algebraically special solutions in higher dimensions," in *Black Holes in Higher Dimensions*, G. Horowitz, ed. Cambridge University Press, Cambridge, UK, Apr., 2011. arXiv:1105.4057.

[28] R. Emparan, T. Harmark, V. Niarchos, and N. A. Obers, "Essentials of Blackfold Dynamics," *JHEP* **03** (2010) 063, arXiv:0910.1601.

[29] R. Emparan, R. Suzuki, and K. Tanabe, "The large D limit of General Relativity," *JHEP* **06** (2013) 009, arXiv:1302.6382.

[30] R. Gregory and R. Laflamme, "Black strings and p-branes are unstable," *Phys. Rev. Lett.* **70** (1993) 2837–2840, arXiv:hep-th/9301052.

[31] L. Lehner and F. Pretorius, "Black Strings, Low Viscosity Fluids, and Violation of Cosmic Censorship," *Phys. Rev. Lett.* **105** (2010) 101102, arXiv:1006.5960.

[32] R. Penrose, "Gravitational collapse: The role of general relativity," *Riv. Nuovo Cim.* **1** (1969) 252–276. [republished in Gen. Rel. Grav. 34, 1141 (2002)].

[33] M. Shibata and H. Yoshino, "Bar-mode instability of rapidly spinning black hole in higher dimensions: Numerical simulation in general relativity," *Phys. Rev.* **D81** (2010) 104035, arXiv:1004.4970.

[34] V. Fock, *The Theory of Space, Time and Gravitation (Second Edition)*. Pergamon, New York, NY, USA, second edition ed., 1964.

[35] G. T. Horowitz and A. Strominger, "Black strings and P-branes," *Nucl. Phys.* **B360** (1991) 197–209.

[36] G. 't Hooft, "A Planar Diagram Theory for Strong Interactions," *Nucl. Phys.* **B72** (1974) 461.

[37] S. Gubser, I. R. Klebanov, and A. M. Polyakov, "Gauge theory correlators from noncritical string theory," *Phys.Lett.* **B428** (1998) 105–114, arXiv:hep-th/9802109.

[38] E. Witten, "Anti-de Sitter space and holography," *Adv.Theor.Math.Phys.* **2** (1998) 253–291, arXiv:hep-th/9802150.

[39] S. de Haro, K. Skenderis, and S. N. Solodukhin, "Gravity in warped compactifications and the holographic stress tensor," *Class. Quant. Grav.* **18** (2001) 3171–3180, arXiv:hep-th/0011230.

[40] S. de Haro, S. N. Solodukhin, and K. Skenderis, "Holographic reconstruction of space-time and renormalization in the AdS / CFT correspondence," *Commun.Math.Phys.* **217** (2001) 595–622, arXiv:hep-th/0002230.

[41] E. Witten, "Anti-de Sitter space, thermal phase transition, and confinement in gauge theories," *Adv.Theor.Math.Phys.* **2** (1998) 505–532, arXiv:hep-th/9803131.

[42] O. Aharony, S. Minwalla, and T. Wiseman, "Plasma-balls in large N gauge theories and localized black holes," *Class.Quant.Grav.* **23** (2006) 2171–2210, arXiv:hep-th/0507219.

[43] S. Hawking, C. Hunter, and M. Taylor, "Rotation and the AdS / CFT correspondence," *Phys.Rev.* **D59** (1999) 064005, arXiv:hep-th/9811056.

[44] T. Wiseman, "Numerical construction of static and stationary black holes," in *Black Holes in Higher Dimensions*, G. Horowitz, ed. Cambridge University Press, Cambridge, UK, Apr., 2012. arXiv:1107.5513.

[45] M. Headrick, S. Kitchen, and T. Wiseman, "A New approach to static numerical relativity, and its application to Kaluza-Klein black holes," *Class.Quant.Grav.* **27** (2010) 035002, arXiv:0905.1822.

[46] A. Adam, S. Kitchen, and T. Wiseman, "A numerical approach to finding general stationary vacuum black holes," *Class.Quant.Grav.* **29** (2012) 165002, arXiv:1105.6347.

[47] S. Hollands, A. Ishibashi, and R. M. Wald, "A Higher dimensional stationary rotating black hole must be axisymmetric," *Commun. Math. Phys.* **271** (2007) 699–722, arXiv:gr-qc/0605106.

[48] P. Figueras, J. Lucietti, and T. Wiseman, "Ricci solitons, Ricci flow, and strongly coupled CFT in the Schwarzschild Unruh or Boulware vacua," *Class.Quant.Grav.* **28** (2011) 215018, arXiv:1104.4489.

[49] P. Figueras and T. Wiseman, "On the existence of stationary Ricci solitons," arXiv:1610.06178 [gr-qc].

[50] G. T. Horowitz and J. E. Santos, "Geons and the Instability of Anti-de Sitter Spacetime," *Surveys Diff. Geom.* **20** (2015) 321–335, arXiv:1408.5906.

[51] O. J. C. Dias, J. E. Santos, and B. Way, "Black holes with a single Killing vector field: black resonators," *JHEP* **12** (2015) 171, arXiv:1505.04793.

[52] P. Figueras and T. Wiseman, "Stationary holographic plasma quenches and numerical methods for non-Killing horizons," *Phys. Rev. Lett.* **110** (2013) 171602, arXiv:1212.4498.

[53] S. Fischetti, D. Marolf, and J. Santos, "AdS flowing black funnels: Stationary AdS black holes with non-Killing horizons and heat transport in the dual CFT," *Class.Quant.Grav.* **30** (2013) 075001, arXiv:1212.4820.

[54] D. M. DeTurck, "Deforming metrics in the direction of their ricci tensors," *J. Differential Geom.* **18** no. 1, (1983) 157–162. http://projecteuclid.org/euclid.jdg/1214509286.

[55] P. Figueras and S. Tunyasuvunakool, "CFTs in rotating black hole backgrounds," *Class. Quant. Grav.* **30** (2013) 125015, arXiv:1304.1162.

[56] S. Fischetti and J. E. Santos, "Rotating Black Droplet," *JHEP* **07** (2013) 156, arXiv:1304.1156.

[57] S. Hawking, "Particle Creation by Black Holes," *Commun.Math.Phys.* **43** (1975) 199–220.

[58] N. Birrell and P. Davies, "Quantum fields in curved space," *Cambridge University Press* (1982) 340p.

[59] V. E. Hubeny, D. Marolf, and M. Rangamani, "Hawking radiation in large N strongly-coupled field theories," *Class.Quant.Grav.* **27** (2010) 095015, arXiv:0908.2270.

[60] V. E. Hubeny, D. Marolf, and M. Rangamani, "Black funnels and droplets from the AdS C-metrics," *Class.Quant.Grav.* **27** (2010) 025001, arXiv:0909.0005.

[61] V. E. Hubeny, D. Marolf, and M. Rangamani, "Hawking radiation from AdS black holes," *Class.Quant.Grav.* **27** (2010) 095018, arXiv:0911.4144.

[62] J. E. Santos and B. Way, "Black Funnels," *JHEP* **1212** (2012) 060, arXiv:1208.6291.

[63] B. S. Kay and R. M. Wald, "Theorems on the Uniqueness and Thermal Properties of Stationary, Nonsingular, Quasifree States on Space-Times with a Bifurcate Killing Horizon," *Phys.Rept.* **207** (1991) 49–136.

[64] P. Figueras and T. Wiseman, "Gravity and large black holes in Randall-Sundrum II braneworlds," *Phys.Rev.Lett.* **107** (2011) 081101, arXiv:1105.2558.

[65] L. Randall and R. Sundrum, "A Large mass hierarchy from a small extra dimension," *Phys.Rev.Lett.* **83** (1999) 3370–3373, arXiv:hep-ph/9905221.

[66] L. Randall and R. Sundrum, "An Alternative to compactification," *Phys.Rev.Lett.* **83** (1999) 4690–4693, arXiv:hep-th/9906064.

[67] T. Tanaka, "Classical black hole evaporation in Randall-Sundrum infinite brane world," *Prog.Theor.Phys.Suppl.* **148** (2003) 307–316, arXiv:gr-qc/0203082.

[68] R. Emparan, A. Fabbri, and N. Kaloper, "Quantum black holes as holograms in AdS brane worlds," *JHEP* **0208** (2002) 043, arXiv:hep-th/0206155.

[69] R. Emparan, J. Garcia-Bellido, and N. Kaloper, "Black hole astrophysics in AdS brane worlds," *JHEP* **0301** (2003) 079, arXiv:hep-th/0212132.

[70] J. E. Santos and B. Way, "Black Droplets," *JHEP* **08** (2014) 072, arXiv:1405.2078.

[71] S. Fischetti and D. Marolf, "Flowing funnels: heat sources for field theories and the ads 3 dual of cft 2 hawking radiation," *Class.Quant.Grav.* **29** (2012) 105004, arXiv:1202.5069.

[72] J. Garriga and T. Tanaka, "Gravity in the brane world," *Phys.Rev.Lett.* **84** (2000) 2778–2781, arXiv:hep-th/9911055.

[73] O. J. Dias, P. Figueras, R. Monteiro, H. S. Reall, and J. E. Santos, "An instability of higher-dimensional rotating black holes," *JHEP* **1005** (2010) 076, arXiv:1001.4527.

[74] M. T. Anderson, P. T. Chrusciel, and E. Delay, "Nontrivial, static, geodesically complete, vacuum space-times with a negative cosmological constant," *JHEP* **0210** (2002) 063, arXiv:gr-qc/0211006.

[75] M. Headrick and T. Wiseman, "Ricci flow and black holes," *Class.Quant.Grav.* **23** (2006) 6683–6708, arXiv:hep-th/0606086.

[76] D. Gross, M. Perry, and L. Yaffe, "Instability of Flat Space at Finite Temperature," *Phys.Rev.* **D25** (1982) 330–355.

[77] R. Monteiro, M. J. Perry, and J. E. Santos, "Semiclassical instabilities of Kerr-AdS black holes," *Phys.Rev.* **D81** (2010) 024001, arXiv:0905.2334.

[78] R. Gregory, "Braneworld black holes," *Lect. Notes Phys.* **769** (2009) 259–298, arXiv:0804.2595.

[79] M. Cvetic, H. Lu, and C. Pope, "Gauged six-dimensional supergravity from massive type IIA," *Phys.Rev.Lett.* **83** (1999) 5226–5229, arXiv:hep-th/9906221.

[80] A. Brandhuber and Y. Oz, "The D-4 - D-8 brane system and five-dimensional fixed points," *Phys.Lett.* **B460** (1999) 307–312, arXiv:hep-th/9905148.

[81] N. Tanahashi and T. Tanaka, "Black holes in braneworld models," *Prog.Theor.Phys.Suppl.* **189** (2011) 227–268, arXiv:1105.2997.

[82] H. Casimir, "On the Attraction Between Two Perfectly Conducting Plates," *Indag.Math.* **10** (1948) 261–263.

[83] P. C. W. Davies, "Quantum vacuum stress without regularization in two-dimensional space-time," *Proc.Roy.Soc.Lond. A.* **354** (1977) 529–532.

[84] S. Christensen and S. Fulling, "Trace Anomalies and the Hawking Effect," *Phys.Rev.* **D15** (1977) 2088–2104.

[85] C. J. Fewster, "Lectures on quantum energy inequalities," arXiv:1208.5399.

[86] C. J. Fewster and M. J. Pfenning, "Quantum energy inequalities and local covariance. I. Globally hyperbolic spacetimes," *J.Math.Phys.* **47** (2006) 082303, arXiv:math-ph/0602042.

[87] K. Howard and P. Candelas, "Quantum stress tensor in Schwarzschild space-time," *Phys.Rev.Lett.* **53** (1984) 403–406.

[88] P. Figueras and S. Tunyasuvunakool, "Localized Plasma Balls," *JHEP* **06** (2014) 025, arXiv:1404.0018.

[89] S. Hawking and D. N. Page, "Thermodynamics of Black Holes in anti-De Sitter Space," *Commun.Math.Phys.* **87** (1983) 577.

[90] D. Marolf, M. Rangamani, and T. Wiseman, "Holographic thermal field theory on curved spacetimes," *Class. Quant. Grav.* **31** (2014) 063001, arXiv:1312.0612.

[91] G. T. Horowitz and R. C. Myers, "The AdS / CFT correspondence and a new positive energy conjecture for general relativity," *Phys.Rev.* **D59** (1998) 026005, arXiv:hep-th/9808079.

[92] R. Emparan and G. Milanesi, "Exact Gravitational Dual of a Plasma Ball," *JHEP* **0908** (2009) 012, arXiv:0905.4590.

[93] J. Plebanski and M. Demianski, "Rotating, charged, and uniformly accelerating mass in general relativity," *Annals Phys.* **98** (1976) 98–127.

[94] R. Emparan, G. T. Horowitz, and R. C. Myers, "Exact description of black holes on branes," *JHEP* **0001** (2000) 007, arXiv:hep-th/9911043.

[95] J. M. Maldacena, "Wilson loops in large N field theories," *Phys.Rev.Lett.* **80** (1998) 4859–4862, arXiv:hep-th/9803002.

[96] S.-J. Rey and J.-T. Yee, "Macroscopic strings as heavy quarks in large N gauge theory and anti-de Sitter supergravity," *Eur.Phys.J.* **C22** (2001) 379–394, arXiv:hep-th/9803001.

[97] P. Figueras and S. Tunyasuvunakool, "Black rings in global anti-de Sitter space," *JHEP* **03** (2015) 149, arXiv:1412.5680.

[98] R. Emparan and H. S. Reall, "Black Rings," *Class.Quant.Grav.* **23** (2006) R169, arXiv:hep-th/0608012.

[99] B. Kleihaus, J. Kunz, and E. Radu, "Black ringoids: spinning balanced black objects in $d \geq 5$ dimensions – the codimension-two case," *JHEP* **01** (2015) 117, arXiv:1410.0581.

[100] R. Emparan, T. Harmark, V. Niarchos, and N. A. Obers, "New Horizons for Black Holes and Branes," *JHEP* **1004** (2010) 046, arXiv:0912.2352.

[101] H. K. Kunduri and J. Lucietti, "An infinite class of extremal horizons in higher dimensions," *Commun.Math.Phys.* **303** (2011) 31–71, arXiv:1002.4656.

[102] O. J. C. Dias, J. E. Santos, and B. Way, "Rings, Ripples, and Rotation: Connecting Black Holes to Black Rings," *JHEP* **07** (2014) 045, arXiv:1402.6345.

[103] R. Emparan, P. Figueras, and M. Martinez, "Bumpy black holes," *JHEP* **12** (2014) 072, arXiv:1410.4764.

[104] R. Emparan and H. S. Reall, "Black Holes in Higher Dimensions," *Living Rev.Rel.* **11** (2008) 6, arXiv:0801.3471.

[105] S. Bhattacharyya, S. Lahiri, R. Loganayagam, and S. Minwalla, "Large rotating AdS black holes from fluid mechanics," *JHEP* **0809** (2008) 054, arXiv:0708.1770.

[106] B. Carter, "Hamilton-Jacobi and Schrodinger separable solutions of Einstein's equations," *Commun.Math.Phys.* **10** (1968) 280.

[107] G. Gibbons, H. Lu, D. N. Page, and C. Pope, "The General Kerr-de Sitter metrics in all dimensions," *J.Geom.Phys.* **53** (2005) 49–73, arXiv:hep-th/0404008.

[108] S. Lahiri and S. Minwalla, "Plasmarings as dual black rings," *JHEP* **0805** (2008) 001, arXiv:0705.3404.

[109] R. Emparan and H. S. Reall, "Generalized Weyl solutions," *Phys.Rev.* **D65** (2002) 084025, arXiv:hep-th/0110258.

[110] T. Harmark, "Stationary and axisymmetric solutions of higher-dimensional general relativity," *Phys.Rev.* **D70** (2004) 124002, arXiv:hep-th/0408141.

[111] M. M. Caldarelli, R. Emparan, and M. J. Rodriguez, "Black Rings in (Anti)-deSitter space," *JHEP* **0811** (2008) 011, arXiv:0806.1954.

[112] R. Emparan, "Rotating circular strings, and infinite nonuniqueness of black rings," *JHEP* **0403** (2004) 064, arXiv:hep-th/0402149.

[113] O. J. Dias and J. P. Lemos, "Pair of accelerated black holes in anti-de Sitter background: AdS C metric," *Phys.Rev.* **D67** (2003) 064001, arXiv:hep-th/0210065.

[114] P. R. Amestoy, I. S. Duff, and J.-Y. L'Excellent, "Multifrontal parallel distributed symmetric and unsymmetric solvers," *Comput. Meth. Appl. Mech. Eng.* **184** no. 2, (2000) 501–520.

[115] P. Amestoy, A. Buttary, A. Guermouche, J.-Y. L'Excellent, and B. Ucar, "MUMPS Users' guide," 2014. http://mumps.enseeiht.fr/.

[116] O. Schenk and K. Gärtner, "Solving unsymmetric sparse systems of linear equations with pardiso," *Future Generation Computer Systems* **20** no. 3, (2004) 475–487.

[117] Intel Corporation, "Reference Manual for Intel MKL PARDISO," 2014. https://software.intel.com/en-us/node/521677.

[118] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, "Efficient management of parallelism in object oriented numerical software libraries," in *Modern Software Tools in*

Scientific Computing, E. Arge, A. M. Bruaset, and H. P. Langtangen, eds., pp. 163–202. Birkhäuser Press, 1997.

[119] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, K. Rupp, B. F. Smith, and H. Zhang, "PETSc users manual," Tech. Rep. ANL-95/11 - Revision 3.5, Argonne National Laboratory, 2014. http://www.mcs.anl.gov/petsc.

[120] A. Magnon, "On Komar integrals in asymptotically anti-de Sitter space-times," *J.Math.Phys.* **26** (1985) 3112–3117.

[121] P. T. Chrusciel, D. Maerten, and P. Tod, "Rigid upper bounds for the angular momentum and centre of mass of non-singular asymptotically anti-de Sitter space-times," *JHEP* **0611** (2006) 084, arXiv:gr-qc/0606064.

[122] H. Elvang, R. Emparan, and A. Virmani, "Dynamics and stability of black rings," *JHEP* **0612** (2006) 074, arXiv:hep-th/0608076.

[123] G. Gibbons, M. Perry, and C. Pope, "The First law of thermodynamics for Kerr-anti-de Sitter black holes," *Class.Quant.Grav.* **22** (2005) 1503–1526, arXiv:hep-th/0408217.

[124] S. Hawking and H. Reall, "Charged and rotating AdS black holes and their CFT duals," *Phys.Rev.* **D61** (2000) 024014, arXiv:hep-th/9908109.

[125] V. Cardoso, Ó. J. Dias, G. S. Hartnett, L. Lehner, and J. E. Santos, "Holographic thermalization, quasinormal modes and superradiance in Kerr-AdS," *JHEP* **1404** (2014) 183, arXiv:1312.5323.

[126] R. Monteiro, M. J. Perry, and J. E. Santos, "Thermodynamic instability of rotating black holes," *Phys.Rev.* **D80** (2009) 024041, arXiv:0903.3256.

[127] B. F. Whiting, "Mode Stability of the Kerr Black Hole," *J. Math. Phys.* **30** (1989) 1301.

[128] M. Dafermos and I. Rodnianski, "Decay for solutions of the wave equation on Kerr exterior spacetimes I-II: The cases |a| << M or axisymmetry," arXiv:1010.5132.

[129] M. Dafermos, I. Rodnianski, and Y. Shlapentokh-Rothman, "Decay for solutions of the wave equation on Kerr exterior spacetimes III: The full subextremal case |a| < M," *Ann. Math.* **183** no. 3, (2016) 787–913, arXiv:1402.7034.

[130] M. Dafermos, I. Rodnianski, and Y. Shlapentokh-Rothman, "A scattering theory for the wave equation on Kerr black hole exteriors," arXiv:1412.8379.

[131] L. Lehner and F. Pretorius, "Numerical Relativity and Astrophysics," *Ann. Rev. Astron. Astrophys.* **52** (2014) 661–694, arXiv:1405.4840 [astro-ph.HE].

[132] M. Zilhão, V. Cardoso, C. Herdeiro, L. Lehner, and U. Sperhake, "Testing the nonlinear stability of Kerr-Newman black holes," *Phys. Rev.* **D90** no. 12, (2014) 124088, arXiv:1410.0694 [gr-qc].

[133] D. Christodoulou and S. Klainerman, *The Global nonlinear stability of the Minkowski space*. Princeton University Press, Princeton, NJ, USA, 1993.

[134] M. Dafermos, G. Holzegel, and I. Rodnianski, "The linear stability of the Schwarzschild solution to gravitational perturbations," arXiv:1601.06467.

[135] J. L. Hovdebo and R. C. Myers, "Black rings, boosted strings and Gregory-Laflamme," *Phys. Rev.* **D73** (2006) 084013, arXiv:hep-th/0601079.

[136] J. E. Santos and B. Way, "Neutral Black Rings in Five Dimensions are Unstable," *Phys. Rev. Lett.* **114** (2015) 221101, arXiv:1503.00721.

[137] R. Emparan and R. C. Myers, "Instability of ultra-spinning black holes," *JHEP* **09** (2003) 025, arXiv:hep-th/0308056.

[138] O. J. C. Dias, P. Figueras, R. Monteiro, J. E. Santos, and R. Emparan, "Instability and new phases of higher-dimensional rotating black holes," *Phys. Rev.* **D80** (2009) 111701, arXiv:0907.2248.

[139] M. Shibata and H. Yoshino, "Nonaxisymmetric instability of rapidly rotating black hole in five dimensions," *Phys. Rev.* **D81** (2010) 021501, arXiv:0912.3606.

[140] O. J. C. Dias, G. S. Hartnett, and J. E. Santos, "Quasinormal modes of asymptotically flat rotating black holes," *Class. Quant. Grav.* **31** no. 24, (2014) 245011, arXiv:1402.7047.

[141] R. Emparan, R. Suzuki, and K. Tanabe, "Instability of rotating black holes: large D analysis," *JHEP* **06** (2014) 106, arXiv:1402.6215.

[142] R. Emparan, R. Suzuki, and K. Tanabe, "Evolution and End Point of the Black String Instability: Large D Solution," *Phys. Rev. Lett.* **115** no. 9, (2015) 091102, arXiv:1506.06772.

[143] K. Tanabe, "Black rings at large D," *JHEP* **02** (2016) 151, arXiv:1510.02200.

[144] M. Alcubierre, *Introduction to 3+1 numerical relativity*. International series of monographs on physics. Oxford University Press, Oxford, UK, 2008.

[145] M. Shibata and T. Nakamura, "Evolution of three-dimensional gravitational waves: Harmonic slicing case," *Phys. Rev.* **D52** (1995) 5428–5444.

[146] T. W. Baumgarte and S. L. Shapiro, "On the numerical integration of Einstein's field equations," *Phys. Rev.* **D59** (1999) 024007, arXiv:gr-qc/9810065.

[147] C. Gundlach and J. M. Martin-Garcia, "Symmetric hyperbolicity and consistent boundary conditions for second order Einstein equations," *Phys. Rev.* **D70** (2004) 044032, arXiv:gr-qc/0403019.

[148] C. Bona, T. Ledvinka, C. Palenzuela, and M. Zacek, "General covariant evolution formalism for numerical relativity," *Phys. Rev.* **D67** (2003) 104005, arXiv:gr-qc/0302083.

[149] C. Gundlach, J. M. Martin-Garcia, G. Calabrese, and I. Hinder, "Constraint damping in the Z4 formulation and harmonic gauge," *Class. Quant. Grav.* **22** (2005) 3767–3774, arXiv:gr-qc/0504114.

[150] D. Alic, C. Bona-Casas, C. Bona, L. Rezzolla, and C. Palenzuela, "Conformal and covariant formulation of the Z4 system with constraint-violation damping," *Phys. Rev.* **D85** (2012) 064040, arXiv:1106.2254.

[151] S. Bernuzzi and D. Hilditch, "Constraint violation in free evolution schemes: Comparing BSSNOK with a conformal decomposition of Z4," *Phys. Rev.* **D81** (2010) 084003, arXiv:0912.2920.

[152] J. D. Brown, "Puncture Evolution of Schwarzschild Black Holes," *Phys. Rev.* **D77** (2008) 044018, arXiv:0705.1359 [gr-qc].

[153] M. J. Berger and J. Oliger, "Adaptive mesh refinement for hyperbolic partial differential equations," *Journal of Computational Physics* **53** no. 3, (1984) 484 – 512.

[154] M. K. Horn, "Fourth- and fifth-order, scaled rungs–kutta algorithms for treating dense output," *SIAM Journal on Numerical Analysis* **20** no. 3, (1983) 558–568.

[155] E. Hairer, S. P. Nørsett, and G. Wanner, *Solving Ordinary Differential Equations I (2Nd Revised. Ed.): Nonstiff Problems*. Springer-Verlag New York, Inc., New York, NY, USA, 1993.

[156] K. Clough, P. Figueras, H. Finkel, M. Kunesch, E. A. Lim, and S. Tunyasuvunakool, "GRChombo : Numerical Relativity with Adaptive Mesh Refinement," *Class. Quant. Grav.* **32** no. 24, (2015) 245011, arXiv:1503.03436.

[157] M. Adams, P. Colella, D. Graves, J. Johnson, N. Keen, T. Ligocki, D. Martin, P. McCorquodale, D. Modiano, P. Schwartz, T. Sternberg, and B. Van Straalen, "Chombo software package for amr applications - design document," Tech. Rep. LBNL-6616E, Lawrence Berkeley National Laboratory, 2015. http://crd.lbl.gov/assets/pubs_presos/chomboDesign.pdf.

[158] H. Kreiss and J. Oliger, *Methods for the approximate solution of time dependent problems*. No. 10 in GARP publications series. World Meteorological Organization, Geneva, 1973.

[159] A. Lapidus, "A detached shock calculation by second-order finite differences," *Journal of Computational Physics* **2** no. 2, (1967) 154 – 177.

[160] F. Pretorius, "Numerical relativity using a generalized harmonic decomposition," *Class. Quant. Grav.* **22** (2005) 425–452, arXiv:gr-qc/0407110.

[161] M. Zilhao, H. Witek, U. Sperhake, V. Cardoso, L. Gualtieri, C. Herdeiro, and A. Nerozzi, "Numerical relativity for D dimensional axially symmetric space-times: formalism and code tests," *Phys. Rev.* **D81** (2010) 084052, arXiv:1001.2302.

[162] M. Alcubierre, S. Brandt, B. Bruegmann, D. Holz, E. Seidel, R. Takahashi, and J. Thornburg, "Symmetry without symmetry: Numerical simulation of axisymmetric systems using Cartesian grids," *Int. J. Mod. Phys.* **D10** (2001) 273–290, arXiv:gr-qc/9908012.

[163] W. G. Cook, P. Figueras, M. Kunesch, U. Sperhake, and S. Tunyasuvunakool, "Dimensional reduction in numerical relativity: Modified cartoon formalism and regularization," *Int. J. Mod. Phys.* **D25** no. 09, (2016) 1641013, arXiv:1603.00362.

[164] S. Husa, I. Hinder, and C. Lechner, "Kranc: A Mathematica application to generate numerical codes for tensorial evolution equations," *Comput. Phys. Commun.* **174** (2006) 983–1004, arXiv:gr-qc/0404023.

[165] J. Thornburg, "Event and apparent horizon finders for 3+1 numerical relativity," *Living Rev. Rel.* **10** (2007) 3, arXiv:gr-qc/0512169.

[166] B. Fornberg, "Calculation of weights in finite difference formulas," *SIAM Rev* **40** (1998) 685–691.

[167] R. G. Keys, "Cubic convolution interpolation for digital image processing," *IEEE Trans. Acoust., Speech, Signal Process* **29** no. 6, (1981) 1153–1160.

[168] E. H. W. Meijering, K. J. Zuiderveld, and M. A. Viergever, "Image reconstruction by convolution with symmetrical piecewise nth-order polynomial kernels," *IEEE Transactions on Image Processing* **8** (1999) 192–201.

[169] P. Figueras, M. Kunesch, and S. Tunyasuvunakool, "End Point of Black Ring Instabilities and the Weak Cosmic Censorship Conjecture," *Phys. Rev. Lett.* **116** no. 7, (2016) 071102, arXiv:1512.04532.

[170] P. Figueras, K. Murata, and H. S. Reall, "Black hole instabilities and local Penrose inequalities," *Class. Quant. Grav.* **28** (2011) 225030, arXiv:1107.5785.