The Spectral Theory of Multiboundary Wormholes

Thesis by
Shaun Cohn Maguire

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2018
Defended May 8, 2018

Shaun Cohn Maguire
ORCID: 0000-0003-3917-8562

Acknowledgements

Caltech is my Hogwarts. It's impossible to describe how significantly finding this institution has impacted my life. On my first day at Caltech, I felt like Harry Potter transitioning from Muggle life to the world of wizardry. It was the first place where I felt comfortable being myself. It transformed me from self-identifying as an introvert to an extrovert, as I found people that I could speak freely with and learn from. I'm deeply indebted to Caltech and its community. I hope this is only the beginning of my relationship with the institution.

At Caltech I've had a number of exceptional mentors, but especially my advisor John Preskill. John has a preternatural intuition that enables him to dance the line between being too hands-off versus too hands-on. He provides an extremely dynamic research group but minimal direct feedback on his students' work. John has a unique gift of being able to make small but impactful comments at exactly the right moments. Many student-advisor relationships motivate the students through fear. With John, the motivation comes out of respect. John took a risk by letting me into his research group and I'll forever be grateful to him for this. I'm also grateful to the rest of my thesis and candidacy committee members, which included: Sergei Gukov, Alexei Kitaev, Oskar Painter, and Mark Wise.

I had a complicated path at Caltech. I originally matriculated as a PhD in the CDS department with the intention of working with Jerry Marsden. Tragically, he passed away during my first year. I was the only PhD student to enter into that option that year, which was an amazing experience given the quality of its faculty. I received exceptional mentorship from Richard Murray, John Doyle, and especially Houman Owhadi. Houman went to bat for me when I was recruited to work at DARPA and I wouldn't have been allowed back to Caltech without his support.

After Marsden's passing I transferred into the mathematics department, which was my undergraduate major. I was only in that department for a quarter before being recruited to DARPA but during that short time I had a few highly impactful conver-

sations with Nikolai Makarov. He introduced me to the book “Indra’s Pearls” and also to the ideas of glueing Riemann surfaces together (conformal weldings) and to quasiconformal maps. A combined two hours of conversation planted seeds that were grew into the kernel of this thesis.

When I returned from DARPA, I joined Preskill’s research group. I drew a huge amount of inspiration from the intellectual diversity of this group, with interests ranging from quantum error-correction to black holes – and then I got to see these interests merge! Some of the regular attendees during my time in John’s Group included: Ning Bao, Mario Berta, Michael Beverland, Thom Bohdanowicz, Peter Brooks, Todd Brun, Elizabeth Crosson, Nicolas Delfosse, Glen Evenbly, Bill Fefferman, Matthew Fishman, Steve Flammia, Alexey Gorshkov, David Gosset, Jeongwan Haah, Nick Hunter-Jones, Joe Iverson, Stacey Jeffery, Stephen Jordan, Isaac Kim, Alexander Kubica, Olivier Landon-Cardinal, Netanel Lindner, Roger Mong, Evgeny Mozgunov, Kirill Shtengel, Sujeet Shukla, Kristan Temme, Beni Yoshida, Nicole Yunger Halpern, and especially Spiros Michalakis; who became a collaborator, co-conspirator, and close friend.

While at Caltech, I also benefited greatly from the broader IQIM community and from Sean Carroll’s quantum cosmology working group. In the physics community more broadly, I am indebted to the Simons Foundation for funding the “It from Qubit” collaboration. I met my primary collaborators at some of their workshops, especially Xi Dong, Alex Maloney and Henry Maxfield. With Henry in particular, I encountered a smarter and more technically sophisticated intellectual dual – we had been thinking about similar questions but from different vantage points. In particular we were both obsessed with understanding simple examples deeply.

From my time at DARPA, I’m grateful to Regina Dugan and Chris White, both of whom changed my life. From my time at USC I’m grateful to Jiangfeng Zhang and Richard Arratia. From Stanford I’m grateful to Persi Diaconis and Bala Rajaratnam. I’m grateful to my start-up brothers Dmitriy (Tseliakhovich) Starson, Joe Meyerowitz, Matt Kraning, and Tim Junio. And of course I’m grateful to the National Science Foundation (GRFP) and Department of Defense (NDSEG) for the research fellowships – and therefore academic freedom – they provided.

One of the great lessons I’ve learned in my life is the latent power of brief interactions. Some of the most influential events in my life have been brief interactions. Some of the most notable include: Arnold Beckman catalyzing a scientific curiosity through probing questions about the tides or weather as he’d get pushed by my childhood

house in his wheel chair; Nikolai Makarov introducing me to the book “Indra’s Pearls”; a cousin who helped me build my first computer when I was seven; a conversation with Patrick Hayden about multiboundary wormholes; and a childhood friends’ father (Doug Borcomon) letting me glance through his personal telescope. Internalizing this lesson, I vow to pay it forward by never discounting the potential significance of brief moments of mentorship.

Finally, I need to thank my family without whose support there’s no hope I’d be where I am. I’d especially like to thank my soon-to-be wife Liana, my mother Lisa, father Kevin, sister Shannon, grandparents Phil and Amy and everyone else. Infinite love to you all.

Abstract

This thesis is primarily based on the paper (Dong et al., 2018) which we'll refer to in this thesis as [DMMM]. This paper studies phase transitions in three dimensional quantum gravity (3D gravity). The main technical tools are the spectral theory of hyperbolic manifolds, especially as they're formulated in the study of Kleinian groups.

This thesis will work through background material in Chapters 1-3, will summarize the key results of [DMMM] in 4, and then will briefly describe unpublished results in 5.

Published Content and Contributions

Dong, Xi et al. (2018). "Phase transitions in 3D gravity and fractal dimension". In: *Accepted for publication in JHEP; arXiv:1802.07275*.

S.M. independently originated the idea for this paper and was heavily involved in the calculations as well as the writing of the paper.

Table of Contents

List of Illustrations

Number *Page*

List of Tables

Number *Page*

Chapter 1

Introduction

As the media has sensationalized to death, one of the most outstanding questions in modern physics is to **discover** and then **understand** a theory of quantum gravity ("QG"). Physicists have already discovered candidate theories, such as string theory. However, understanding these theories – carrying out all of the relevant computations to confirm that they are consistent with Nature and then doing experiments to verify their novel predictions – is still beyond our ability. Surprisingly, without knowing the specific theory of QG that guides Nature's hand, we're still able to say a number of universal things that must be true for any theory of QG. The most prominent example being the holographic principle which comes from the entropy of black holes being proportional to the surface area encapsulated by the black hole's horizon (a naive guess says the entropy should be proportional to the volume of the black hole; such as the entropy of a glass of water). Universal statements such as this serve as guideposts and consistency checks as we try to understand QG.

It's exceedingly rare to find universal statements that are true in physically realistic models of quantum gravity. The holographic principle is one such example, but it pretty much stands alone in its power and applicability. By physically realistic I mean: (3 + 1)-dimensional and with the curvature of the universe being either flat, or very mildly positively curved. However, we can make additional simplifying assumptions where it's easier to find universal properties. For example, we can reduce the number of spatial dimensions so that we're considering $(2+1)$-dimensional quantum gravity (3D gravity). Or we can investigate spacetimes that are negatively curved (anti-de Sitter space) as in the AdS/CFT correspondence. Or we can make both assumptions which is the setting for this thesis. The hope is that what's learned in these limited situations will back-propagate insights towards reality.

The motivation for going to (2 + 1)-dimensions is that gravity (general relativity) is much simpler here. We'll say more later but in (2 + 1)-dimensions the physical degrees of freedom of Einstein's equations are exactly balanced by the flexibility offered by the symmetries of the theory which means there are no local degrees of freedom. This in turn means there are no gravitiational waves in the context of general relativity and no gravitons in the context of quantum gravity.

The standard motivation for considering negatively curved spacetimes is that it puts us in the domain of AdS/CFT, which is the best understood model of quantum gravity. However, it's worth pointing out that most of the results in this thesis don't rely on AdS/CFT. We consider negatively curved spacetimes (negatively curved Lorentzian manifolds) because they're related to what mathematicians call hyperbolic manifolds (negatively curved Euclidean manifolds), and mathematicians know a great deal about these objects. It's just a helpful coincidence that because we're working with negatively curved manifolds we also get to unpack our statements in AdS/CFT.

This thesis makes both of these assumptions: we'll be studying gravity in negatively curved 3-manifolds. This work was originally motivated out of a desire to deeply understand the simplest non-trivial toy examples in AdS/CFT, which are called multiboundary wormholes (MBWs.) MBWs are given as quotients of AdS_3 by discrete subgroups of Isom(AdS_3). To every Lorentzian MBW there is a corresponding hyperbolic 3-manifold that is obtained by analytic continuation. The goal of this work was to shed light on how things like Ryu-Takayanagi and bulk-boundary reconstruction behave in MBWs by using the mathematical rigidity of hyperbolic 3-manifolds. Progress was made towards better understanding both of these questions, but the main result of this thesis was the discovery of a universal phase transition – similar to the Hawking-Page phase transition between thermal AdS and AdS black holes – which follows from properties of the spectrum of the associated hyperbolic 3-manifold.

Outline

This thesis is organized as follows:

- In Chapter 2, we review multiboundary wormholes, sketch out how a few examples work, and review the quantum gravity partition function.

- In Chapter 3, we provide the necessary mathematics background. There are a number of significant results that aren't found in the physics literature so we spend extra time introducing them as a placeholder for future work.
- In Chapter 4, we review the key results from [DMMM] and provide additional intuition.
- In Chapter 5, we describe some of the non-published results obtained regarding MBWs, especially as related to analytic continuation in AdS_3 and classification of MBWs.
- In Chapter 6, we introduce some of the follow-up questions that we would like to pursue in the future.

Chapter 2

Introduction to Multiboundary Wormholes

This chapter describes multiboundary wormholes and some of their associated physics. First, we recall the definition of AdS_3 (both the Lorentzian and Euclidean versions). Then we introduce multiboundary wormholes as quotients. We then work through a few examples. Finally, we review one general form of the quantum gravity partition function.

2.1 Anti de-Sitter space

Recall that AdS_{d+1} is the $d+1$-dimensional vacuum solution of Einstein's field equations with constant negative curvature. In this thesis it's helpful to differentiate between Lorentzian and Euclidean solutions by adding a superscript (AdS^L_{d+1} / AdS^E_{d+1}). The Lorentzian version can be thought of as an infinite cylinder given by $\mathbb{R} \times \mathbb{H}^d$, where each time $t \in \mathbb{R}$ indexes a copy of d-dimensional hyperbolic space. The Euclidean version is just $d+1$ dimensional hyperbolic space $\mathbb{H}^{d+1}$. In this thesis we'll restrict to $d = 2$.

$\mathbf{AdS^L_3}$

There are a few different coordinate systems that will be useful. This section uses the notation of (Maxfield, 2015) [Maxfield1] because it's particularly well suited for studying quotients and especially their entanglement properties via (Ryu and Takayanagi, 2006). For more detail we recommend referring to [Maxfield1]. The most obvious coordinate system comes from thinking of AdS^L_3 as a submanifold of

$\mathbb{R}^{2,2}$ given by

$$AdS_3^L = \left\{(U,V,X,Y) \in \mathbb{R}^{2,2} \;\middle|\; U^2 + V^2 - X^2 - Y^2 = L_{AdS}^2\right\}. \tag{2.1}$$

The metric on AdS_3^L is the induced metric from that on $\mathbb{R}^{2,2}$

$$ds^2 = -dU^2 - dV^2 + dX^2 + dY^2. \tag{2.2}$$

This presentation can be rewritten showing that AdS_3^L is isomorphic to the group manifold $SL(2,\mathbb{R})$. First, write $\mathbb{R}^{2,2}$ as

$$\mathbb{R}^{2,2} = \left\{p = \begin{pmatrix} U+X & Y-V \\ Y+V & U-X \end{pmatrix} \text{ where } U,V,X,Y \in \mathbb{R}\right\} \tag{2.3}$$

with metric given by

$$ds^2 = -L_{AdS}^2 \det(p). \tag{2.4}$$

The constraint to obtain AdS_3^L as a submanifold of this is simply that $\det(p) = 1$ showing that indeed $AdS_3^L \cong SL(2,\mathbb{R})$. The group manifold presentation will be helpful when we introduce quotients of AdS_3^L.

It's also important to introduce coordinates that make the temporal structure more obvious. One common coordinate system is given by the substitution

$$X = r\cos\phi \tag{2.5}$$

$$Y = r\sin\phi \tag{2.6}$$

$$U = \sqrt{1+r^2}\cos t \tag{2.7}$$

$$V = \sqrt{1+r^2}\sin t \tag{2.8}$$

with induced metric

$$ds^2 = -(1+r^2)dt^2 + \frac{dr^2}{1+r^2} + r^2 d\phi^2 \tag{2.9}$$

where ϕ and t are periodic in 2π. Note that all of these constructions yield closed time-like curves. This is normally addressed by lifting from AdS_3^L to the universal cover. This can be done by extending $t \in S^1$ to $t \in \mathbb{R}$. However, when dealing with MBWs, the closed time-like curves disappear so we don't need to consider this subtlety.

It should be plausible from the submanifold presentation that the the isometry group of AdS_3^L is

$$\text{Isom}\left(AdS_3^L\right) = SO(2,2). \tag{2.10}$$

We can use the group manifold presentation to see that the connected part of $\text{Isom}(AdS_3^L)$ is isomorphic to $(SL(2,\mathbb{R}) \times SL(2,\mathbb{R}))/\mathbb{Z}_2$. There are a few ways to show this but one is to chose a group action, such as $p \mapsto g_L p g_R^T$ for g_L and g_R each in $SL(2,\mathbb{R})$ with the equivalence relation $(g_L, g_R) \sim (-g_L, -g_R)$. In this group action, $g_L, g_R, p \in SL(2,\mathbb{R})$ and the multiplication is matrix multiplication. Note that this choice of action is not the most natural from the perspective of group theory, where the action would normally be written as $p \mapsto g_L p g_R^{-1}$. This action is chosen so that when $g_L = g_R$ the action fixes the $t = 0$ plane. We'll exploit this structure and often abuse notation by writing $\Gamma \subseteq PSL(2,\mathbb{R})$ when in reality we mean elements $g_L = g_R = \Gamma$ so that the action is (Γ, Γ) acts on p by $p \mapsto \Gamma p \Gamma^T$. Points with $t = 0$ correspond to $V = 0$ which in turn corresponds to symmetric matrices in the group manifold picture. It's easy to check that for $p, \Gamma \in SL(2,\mathbb{R})$ with p symmetric the action $p \mapsto \Gamma p \Gamma^T$ yields another symmetric matrix.

There's one final coordinate system we need to introduce that's also useful when introducing quotients of AdS_3. We can map the $t = 0$ slice of AdS_3^L onto the upper half-plane model of $\mathbb{H}^2$ by setting

$$p = \begin{pmatrix} U+X & Y \\ Y & U-X \end{pmatrix} := \frac{1}{\text{Im}(z)} \begin{pmatrix} |z|^2 & \text{Re}(z) \\ \text{Re}(z) & 1 \end{pmatrix} \tag{2.11}$$

where $z = x + iy \in \mathbb{H}^2$ and this is the upper half-plane model so that $x \in \mathbb{R}$, $y \in \mathbb{R}^+$.

It's worth mentioning that, using the group manifold presentation, the distance between two space-like separated points $p, q \in SL(2,\mathbb{R})$ can be calculated using the formula

$$\ell(p,q) = 2\cosh^{-1}\left(\frac{\text{Tr}(p^{-1}q)}{2}\right). \tag{2.12}$$

This formula is especially convenient when calculating entanglement entropies via the Ryu-Takayanagi formula. Note that there's an ambiguity for points on the boundary which was described eloquently in (Maxfield, 2015).

$\mathbf{AdS_3^E}$

We can analytically continue AdS_3^L by sending $Y \mapsto iY$ so that

$$p = \begin{pmatrix} U+X & Y-V \\ Y+V & U-X \end{pmatrix} \mapsto p_E = \begin{pmatrix} U+X & Y-iV \\ Y+iV & U-X \end{pmatrix} \in SL(2,\mathbb{C}). \tag{2.13}$$

We'll abuse notation and usually also refer to p_E as p. It's also convenient to make a connection with the standard presentation of hyperbolic space $\mathbb{H}^3$ by relabeling

$$U \mapsto T \tag{2.14}$$

$$V \mapsto Y \tag{2.15}$$

$$X \mapsto Z \tag{2.16}$$

$$Y \mapsto X \tag{2.17}$$

so that

$$p_E \mapsto \begin{pmatrix} T+Z & X-iY \\ X+iY & T-Z \end{pmatrix} \tag{2.18}$$

where the constraint $\det(p_E) = 1$ implies that $T^2 - X^2 - Y^2 - Z^2 = 1$. Also note that the metric on $\mathbb{R}^{2,2}$ that we started with gets mapped to

$$ds^2 = -dU^2 + dX^2 + dY^2 + dZ^2 \tag{2.19}$$

on $\mathbb{R}^{3,1}$. In the group manifold language this metric is still given as $ds^2 = -\det(dp)$. We obtain the ball model of $\mathbb{H}^3$ from the substitutions

$$X = \frac{2x}{1-r^2}$$

$$Y = \frac{2y}{1-r^2}$$

$$Z = \frac{2z}{1-r^2}$$

$$T = \frac{1+r^2}{1-r^2}$$

so that

$$ds^2 = \frac{4(dx^2+dy^2+dz^2)}{(1-r^2)^2} \tag{2.20}$$

where $r^2 = x^2 + y^2 + z^2$.

For AdS_3^E the connected part of the isometry group is $SO(3,1) \cong PSL(2,\mathbb{C})$, using the action $p \mapsto gp^\dagger g$. The $t = 0$ slice was obfuscated using the double relabeling but it's given by the p_E with $Y = 0$. Again, these correspond to symmetric matrices in $PSL(2,\mathbb{R})$. The $t = 0$ plane is fixed by $\Gamma \in PSL(2,\mathbb{R}) \subset PSL(2,\mathbb{C})$. The fact that there's a matching diagonal copy of $PSL(2,\mathbb{R})$ in both $\mathrm{Isom}(AdS_3^L)$ and also in $\mathrm{Isom}(AdS_3^E)$ becomes important when trying to perform analytic continuation with quotients.

The formula for the distance between points (lengths of geodesics connecting them) is the same as in the Lorentzian case

$$\ell(p_E, q_E) = 2\cosh^{-1}\left(\frac{\mathrm{Tr}(p^{-1}q)}{2}\right). \tag{2.21}$$

Again, there's a subtletly when computing the distance between boundary points because the distance is infinite. A natural regularization procedure relevant to Ryu-Takayanagi calculations was described in (Maxfield, 2015).

2.2 Quotients of Anti de-Sitter Space

Before describing quotients of AdS we'll motivate why this is a reasonable thing to do by introducing a simple theorem from General Relativity.

Quotients in General Relativity

Recall that Einstein's Field Equations (EFE) for General Relativity take the form

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} \tag{2.22}$$

where $R_{\mu\nu}$ is the Ricci curvature tensor (which only depends on local curvature), R is the Ricci scalar which depends on global curvature, $g_{\mu\nu}$ is the metric tensor, Λ is the cosmological constant, and $T_{\mu\nu}$ is the stress-energy tensor. In this thesis we'll be studying vacuum solutions, so $T_{\mu\nu}$ will be set to zero. Hence the EFE can be written as

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} + \Lambda g_{\mu\nu} = 0. \tag{2.23}$$

If we contract each term with the inverse metric tensor $g^{\mu\nu}$ we obtain a formula for R in terms of Λ and the dimension of the spacetimes manifolds we're considering (let's once again consider $d+1$-dimensional spacetimes)

$$R_{\mu\nu}g^{\mu\nu} - \frac{1}{2}Rg_{\mu\nu}g^{\mu\nu} + \Lambda g_{\mu\nu}g^{\mu\nu} = 0 \tag{2.24}$$

$$\implies R - \frac{d+1}{2}R + (d+1)\Lambda = 0 \tag{2.25}$$

$$\implies R = \frac{2(d+1)}{d-1}\Lambda \tag{2.26}$$

which we can plug back into equation (2.23) to obtain

$$R_{\mu\nu} = \frac{2}{d-1}\Lambda g_{\mu\nu}. \tag{2.27}$$

Rewriting the vacuum field equations in this form shows that they only depend on local curvature. This leads to the following important fact

Theorem 1 *If M is a solution of the vacuum EFE and Γ is a discrete subgroup of $Isom(M)$ then the quotient M/Γ is also a solution of the vacuum EFE.*

This follows trivially from formula 2.27 because quotients by isometries don't change either its left or right hand sides (but they'll change the coordinate ranges and global topology). We need to consider discrete subgroups because otherwise we can create singularities and alter the dimension of the spacetime.

As a simple example for intuition, consider $2+1$-dimensional Minkowski spacetime $\mathbb{R}^{2,1}$ which can be thought of as a stack of copies of $\mathbb{R}^2$ indexed by time. The topological structure is isomorphic to $\mathbb{R}\times\mathbb{R}^2$. The isometry group contains translations in space. Let Γ be the cyclic group generated by a single translation, say translation along the x-axis by 10 meters. The quotient $\mathbb{R}^{2,1}/\Gamma$ where $\Gamma = \langle x \sim x + 10\rangle$ yields a cylindrical spacetime where every constant time slice is a cylinder with diameter 10 meters. The topological structure of this spacetime is $\mathbb{R}\times(\mathbb{R}\times S^1)$, where S^1 is the unit circle and the parentheses is just for emphasis.

Quotients of AdS_3

We can now use theorem 1 and the isometry groups from section 2.1

$$\text{Isom}\left(AdS_3^L\right)^+ = SO(2,2)^+ \cong \frac{SL(2,\mathbb{R})\times SL(2,\mathbb{R})}{\mathbb{Z}_2} \tag{2.28}$$

$$\text{Isom}\left(AdS_3^E\right)^+ = SO(3,1)^+ \cong PSL(2,\mathbb{C}). \tag{2.29}$$

to introduce quotients of AdS_3[1].

Definition 1 *A Multiboundary Wormhole (MBW) is a quotient of AdS_3 (Lorentzian or Euclidean) by a discrete subgroup Γ of $Isom(AdS_3)$ (the Lorentzian or Euclidean isometry group, respectively).*

Whether we're considering Lorentzian or Euclidean MBWs should be clear from the context. The diagonal part of 2.28 and the group 2.29 are extremely well studied, with discrete subgroups of $SL(2,\mathbb{R})$ being called Fuchsian groups and discrete subgroups of $PSL(2,\mathbb{C})$ being called Kleinian groups. The physics of MBWs is intimately tied to the mathematics of these groups.

This construction was introduced in the time-symmetric case in a series of papers (D. R. Brill, 1996; Aminneborg, Bengtsson, D. Brill, et al., 1998); and then in

[1] Note that the superscript + indicates we're considering the connected part of a group.

the general case in (Aminneborg, Bengtsson, and Holst, 1999). It's easiest to understand the construction in the time-symmetric ('non-rotating') case when Γ is a subgroup of the diagonal $PSL(2,\mathbb{R})$ acting on the $t = 0$ slice. One can check that for $g \in \Gamma \subset PSL(2,\mathbb{R})$ the action we chose says we want to identify the points $p \sim gpg^T$. After some algebra you'll see that in terms of the upper half-plane coordinates on the $t = 0$ slice (2.11), this identification corresponds to identifying z with its image after fractional linear transformation gz. More specifically, if $g = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in PSL(2,\mathbb{R})$ then z gets identified with $\frac{az+b}{cz+d}$.

In other words, time-symmetric Lorentzian MBWs are given as quotients AdS_3^L/Γ where $\Gamma \in PSL(2,\mathbb{R})$. The $t = 0$ slice will be given by $\mathbb{H}^2/\Gamma$ which is an open Riemann surface. The global topology has a 'tent-like' form that was described in (D. R. Brill, 1996). We'll work through a few examples in section 2.3. In the more general case, when Γ is not a subset of $PSL(2,\mathbb{R})$ the identifications are less intuitive and these correspond to spacetimes with spinning wormholes in the bulk, where one needs to identify every point with all of its translations by the formula: for every $(g_L, g_R) \in (\Gamma_L, \Gamma_R) \subset SL(2,\mathbb{R}) \times SL(2,\mathbb{R})/\mathbb{Z}_2$ and $p \in AdS_3^L \cong SL(2,\mathbb{R})$ we need to identify $p \mapsto g_L p g_R^T$. It's easy to write this down algebraically but it's hard to gain intuition. Analytic continuation is also more subtle for MBWs with rotation in the bulk. This will be described in section 5.1.

The story in the Euclidean picture is similar, where the MBWs are simple to describe when $\Gamma \subset PSL(2,\mathbb{R}) \subset PSL(2,\mathbb{C})$. We can use the same upper-half-plane coordinates for the $t = 0$ plane because again, using the coordinates from 2.13, the points in the $t = 0$ plane correspond to symmetric matrices in $PSL(2,\mathbb{R})$. This convenience also makes it easy to analytically continue time-symmetric MBWs (in both directions between Lorentzian and Euclidean).

2.3 Examples

The machinery behind these constructions becomes more apparent after investigating a few examples. We'll go through the nitty gritty in simple examples so that someone would be able to reproduce the important calculations if they wanted to (such as doing Ryu-Takayanagi calculations), but one doesn't need to understand these calculations in order to appreciate the high level results in the rest of the paper. If a reader wants to learn more about entanglement entropy in these spacetimes we recommend reading (Maxfield, 2015). If one wants to understand the causal

structures they should consult (D. R. Brill, 1996; Aminneborg, Bengtsson, D. Brill, et al., 1998; D. Brill, 1998).

Static BTZ black hole

We'll start with the simplest example, which is a time-symmetric spacetime obtained as the quotient by a single generator. This is called the static BTZ black hole. It's obtained by the quotient $AdS_3^L/\langle A\rangle$ where A is of the form

$$A = \begin{pmatrix} e^{\frac{r_+}{2}} & 0 \\ 0 & e^{-\frac{r_+}{2}} \end{pmatrix}. \tag{2.30}$$

The notation $\langle G_1, \ldots, G_n\rangle$ for arbitrary generators $G_1, \ldots, G_n$ is meant to denote the associated free group. When there's only one generator, such as in the BTZ case, this notation means the cyclic group generated by A.

We introduced many equivalent coordinate systems above. On the most fundamental level we should think of a quotient by A meaning that to every $p \in SL(2,\mathbb{R})$ we identify all the points $p \sim ApA^T$, $p \sim A^2p(A^2)^T$, etc. It's not hard to trace through what this means but it's tedious. It's much simpler to understand what's happening using the coordinates $z \in \mathbb{H}^2$ that we introduced by identifying the $t = 0$ slice of AdS_3^L with the upper half-plane model of hyperbolic space. Recall that the identification was that

$$p = \begin{pmatrix} U+X & Y \\ Y & U-X \end{pmatrix} := \frac{1}{\mathrm{Im}(z)} \begin{pmatrix} |z|^2 & \mathrm{Re}(z) \\ \mathrm{Re}(z) & 1 \end{pmatrix} \tag{2.31}$$

where $z \in \mathbb{H}^2$ and the action of an element A of the isometry group $p \mapsto ApA^T$ becomes $z \mapsto A.z$ where $A.z$ denotes fractional linear transformation. Hence for this specific example, in the $t = 0$ plane the point z is identified with

$$z \sim A.z = \frac{e^{r_+/2} \cdot z + 0}{0 \cdot z + e^{-r_+/2}} = e^{r_+} z \tag{2.32}$$

Because we're quotienting by the cyclic group $\langle A\rangle$ we need to consider higher powers of A. It's easy to show that for A^n the identification will be

$$z \sim e^{nr_+} z, \tag{2.33}$$

where $n \in \mathbb{Z}$.

Figure 2.1 puts this algebra into picture form. The map $z \mapsto Az$ maps the smaller solid black circle onto the larger one. The dashed red line is a horizon, meaning that

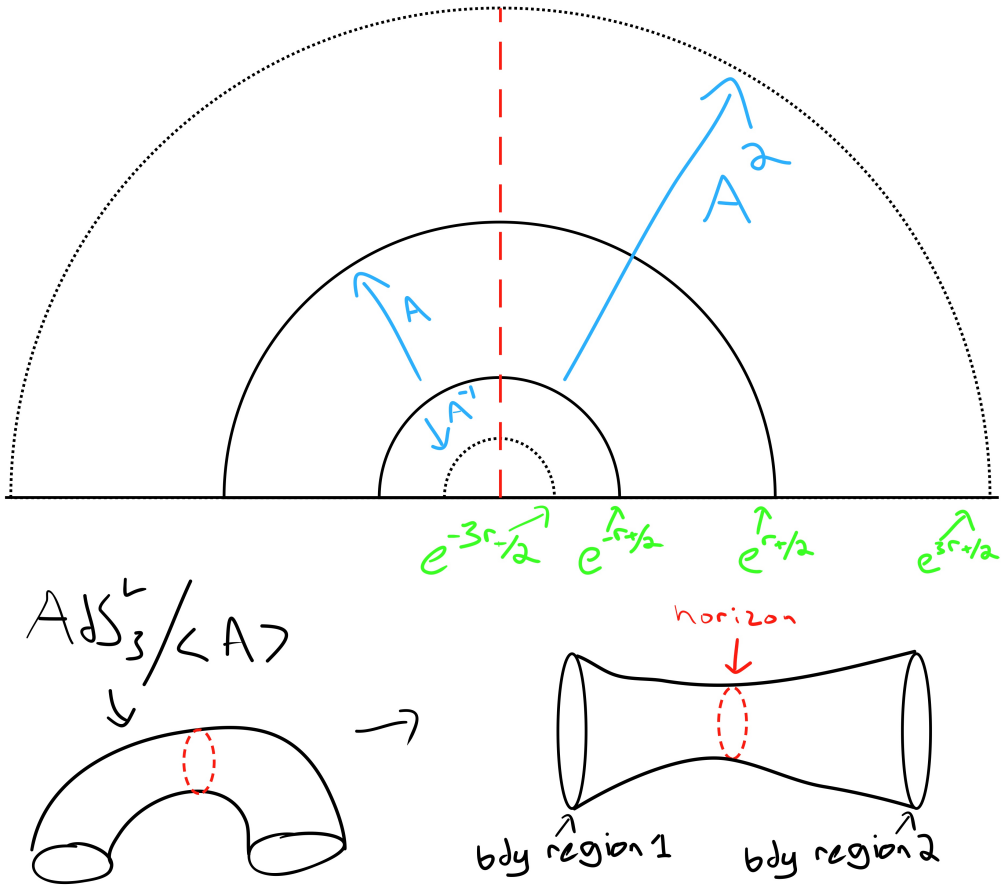

Figure 2.1: This figure shows the $t = 0$ slice of a BTZ black hole. The solid black lines are identified with each other via the generator A. The dashed black lines are examples of higher order identifications. The identification gives a surface that's topologically a cylinder. It has two asymptotic boundary regions which in the original picture live on the x-axis. The dashed red line becomes an event horizon in the quotient. The length of the horizon is r_+. This is only the $t = 0$ slice. The full spacetime diagram is given as a stack of these surfaces index by time. However, the causal structure is a bit subtle. It was described in (D. R. Brill, 1996)

the two asymptotic boundary regions are causally disconnected when one lifts this $t = 0$ slice to a full 2 + 1-dimensional spacetime diagram. See (D. R. Brill, 1996) for the details on this. However, it's worth mentioning that in this form, it's easy to calculate the length of the horizon using the formula for the distance between two points $z_1, z_2 \in \mathbb{H}^2$ in the hyperbolic upper half plane

$$d(z_1, z_2) = \cosh^{-1}\left(1 + \frac{(\mathrm{Re}(z_2) - \mathrm{Re}(z_1))^2 + (\mathrm{Im}(z_2) - \mathrm{Im}(z_1))^2}{2\,\mathrm{Im}(z_1)\,\mathrm{Im}(z_2)}\right) \tag{2.34}$$

so that plugging in the points $z_1 = ie^{-r_+/2}$ and $z_2 = ie^{r_+/2}$ we see that the generator was chosen so that the horizon length is r_+. This can also be seen from equation 2.12 taking p and $q = A.p$ for any arbitrary p corresponding to a point in the $t = 0$ slice. The distance between p and q becomes

$$\ell(p, Ap) = 2\cosh^{-1}\left(\frac{\mathrm{Tr}(p^{-1}q)}{2}\right) \tag{2.35}$$

$$= 2\cosh^{-1}\left(\frac{\mathrm{Tr}(A)}{2}\right) \tag{2.36}$$

$$= 2\cosh^{-1}\left(\frac{e^{r_+/2} + e^{-r_+/2}}{2}\right) \tag{2.37}$$

$$= r_+ \tag{2.38}$$

which for more complicated generators provides a convenient way to calculate horizon lengths.

Rotating BTZ black hole

The simplest example of a non-static multiboundary wormhole is obtained by taking two different cyclic groups living in theft and right copies of $SL(2, \mathbb{R})$ respectively. We'll only provide a sketch for how this example works. Recall that

$$\mathrm{Isom}(AdS_3^L) \cong \frac{SL(2, \mathbb{R}) \times SL(2, \mathbb{R})}{\mathbb{Z}_2}. \tag{2.39}$$

Consider the subgroup $\langle A_L \rangle \times \langle A_R \rangle \subset SL(2, \mathbb{R}) \times SL(2, \mathbb{R})$ where

$$A_L = \begin{pmatrix} e^{\frac{r_+ - r_-}{2}} & 0 \\ 0 & e^{-\frac{r_+ - r_-}{2}} \end{pmatrix} \tag{2.40}$$

$$A_R = \begin{pmatrix} e^{\frac{r_+ + r_-}{2}} & 0 \\ 0 & e^{-\frac{r_+ + r_-}{2}} \end{pmatrix} \tag{2.41}$$

$$\tag{2.42}$$

with $r_+ \geq r_- \geq 0$. The case $r_- = 0$ recovers the static example from the previous subsection. The case $r_+ = r_-$ is the extremal rotating BTZ black hole. In terms of the inner and outer horizon radii (r_- and r_+ respectively) the physical parameters are $M = r_+^2 + r_-^2$ with angular momentum $J = 2r_+r_-$. Again see (Maxfield, 2015) for details on the modern notation and (Aminneborg, Bengtsson, D. Brill, et al., 1998) for details on the causal structure. The reason for including this example is to show that this framework can also incorporate non-static spacetimes.

Three boundary wormhole

We saw in the previous subsections that when there's only one generator we obtain the BTZ black hole (obtained by glueing two circles to each other). When there are two generators there are four initial circles and only two topologically distinct ways of glueing them together. If the initial circles are adjacent then we obtain the three boundary wormhole, see figure 2.2. If the initial circles are opposite each other then we obtain the torus wormhole, see figure 2.3. We'll start with the three boundary wormhole example.

We're only going to introduce the static three boundary wormhole, which corresponds to taking generators that live in the diagonal copy of $SL(2,\mathbb{R})$ inside the full isometry group of AdS_3^L. Consider the generators

$$A = \begin{pmatrix} \cosh\frac{L_1}{2} & \sinh\frac{L_1}{2} \\ \sinh\frac{L_1}{2} & \cosh\frac{L_1}{2} \end{pmatrix} \tag{2.43}$$

$$B = \begin{pmatrix} \cosh\frac{L_2}{2} & e^{\alpha}\sinh\frac{L_2}{2} \\ e^{-\alpha}\sinh\frac{L_2}{2} & \cosh\frac{L_2}{2} \end{pmatrix} \tag{2.44}$$

with $L_1, L_2 > 0$ and α large enough so that $e^{\alpha} > \coth\left(\frac{L_1}{4}\right)\coth\left(\frac{L_2}{4}\right)$. This constraint is chosen so that the initial circles don't overlap, which would lead to a non-physical quotient. The quotient by $\langle A, B\rangle$ gives a spacetime where each constant time slice is a pair of pants geometry. Each funnel is causally disconnected from the others by a horizon. The horizon lengths will be L_1, L_2 and L_3 where L_3 has a complicated form. It's obtained by calculating

$$L_3 = 2\cosh^{-1}\left(\frac{\mathrm{Tr}(-AB^{-1})}{2}\right). \tag{2.45}$$

The matrix $-AB^{-1}$ was chosen because if one traces the path of a point on the third horizon after applying this sequence of matrices they'll return to where they started.

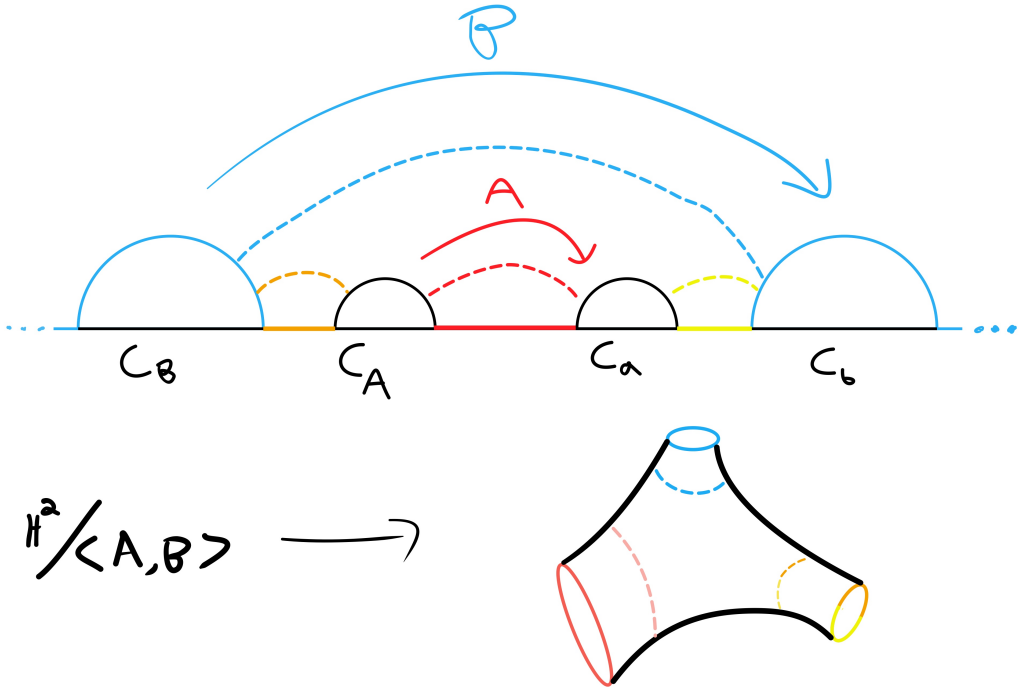

Figure 2.2: This figure shows the $t = 0$ slice of a three boundary wormhole. Note that the identifications are drawn so that everything sits on the upper half plane. In the upper half plane the algebra is simpler but the illustrations are simpler in the disk model. It's helpful to think of the x-axis as being identified at $\pm\infty$. When the blue circles are glued together the two blue boundary segments and the dots are glued together to create a single segment The colors are only there to help one keep track of where different regions go after making identifications. The generator A glues the two black semicircles together. After this identification the solid red segment becomes an asymptotic boundary region. The dashed red line becomes a horizon. The generator B glues the green semicircles together. The solid blue segment in between the green semicircles becomes an asymptotic boundary region. The dashed blue line becomes a horizon. However, there's a third boundary region that's created. It's not as clear from the figure what it corresponds to but it's obtained by gluing the yellow and orange segments together. The yellow and orange dashed arcs get glued together to become the third horizon.

Once again, (Maxfield, 2015) works through this example carefully. The circles $C_A \mapsto C_a$ and $C_B \mapsto C_b$ that are glued together in figure 2.2 have centers and radii (center, radii)

$$C_A = \left(-\frac{\cosh\left(\frac{L_1}{2}\right)}{\sinh\left(\frac{L_1}{2}\right)}, \frac{1}{\sinh\left(\frac{L_1}{2}\right)} \right) \tag{2.46}$$

$$C_a = \left(\frac{\cosh\left(\frac{L_1}{2}\right)}{\sinh\left(\frac{L_1}{2}\right)}, \frac{1}{\sinh\left(\frac{L_1}{2}\right)} \right) \tag{2.47}$$

$$C_B = \left(-\frac{e^{\alpha}\cosh\left(\frac{L_2}{2}\right)}{\sinh\left(\frac{L_2}{2}\right)}, \frac{e^{\alpha}}{\sinh\left(\frac{L_2}{2}\right)} \right) \tag{2.48}$$

$$C_b = \left(\frac{e^{\alpha}\cosh\left(\frac{L_2}{2}\right)}{\sinh\left(\frac{L_2}{2}\right)}, \frac{e^{\alpha}}{\sinh\left(\frac{L_2}{2}\right)} \right) \tag{2.49}$$

One can obtain a spinning three boundary wormhole by considering a more general group $\langle A_L, B_L \rangle \times \langle A_R, B_R \rangle \subset SL(2,\mathbb{R}) \times SL(2,\mathbb{R})$. Note that there will be a variety of constraints on the generators to ensure that the spacetime they yield is physical (which in this case means there won't be any 'bad' singularities).

Torus wormhole

The other topological possibility with two generators is to create the torus wormhole. This happens when A and B pair circles that are across from each other, such as in figure 2.3. A convenient parameterization sets

$$A = \begin{pmatrix} e^{\lambda} & 0 \\ 0 & e^{-\lambda} \end{pmatrix} \tag{2.50}$$

$$B = \begin{pmatrix} \cosh\mu + \cos\alpha \sinh\mu & \sin\alpha \sinh\mu \\ \sin\alpha \sinh\mu & \cosh\mu - \cos\alpha \sinh\mu \end{pmatrix}. \tag{2.51}$$

We can calculate the horizon length by applying formula 2.12 to $ABA^{-1}B^{-1}$, or any of its cyclic permutations. Once again, these words were chosen so that if one traces the path of a point on the horizon after applying this sequence of letters then one ends up back where one started.

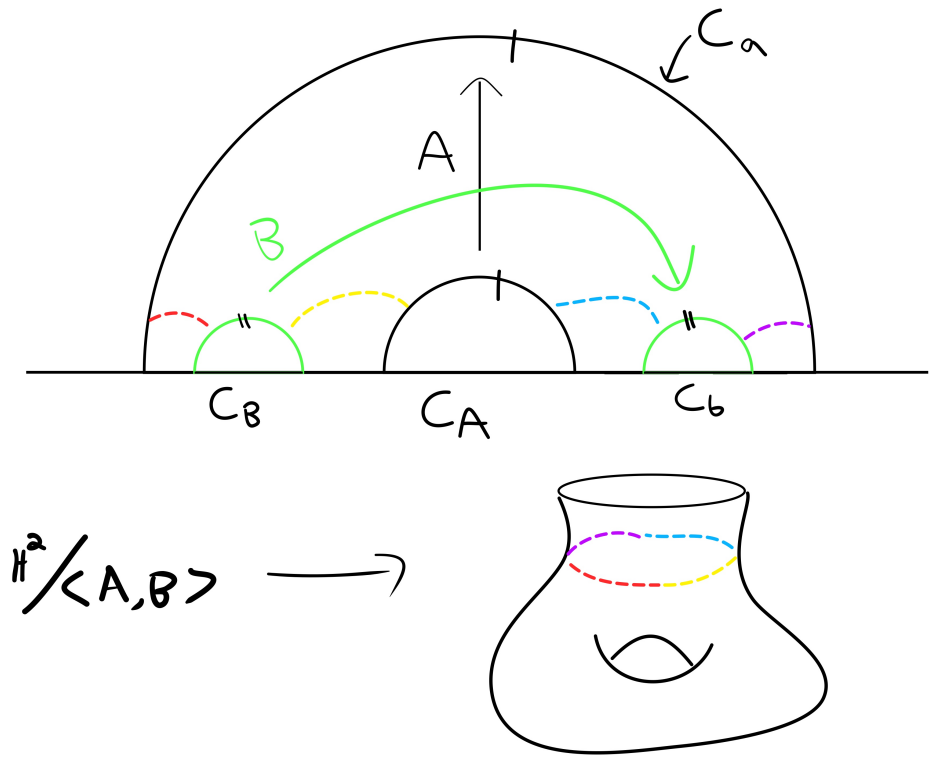

Figure 2.3: This figure shows the $t = 0$ slice of a torus wormhole. Circle C_A is glued to C_a. Circle C_B is glued to C_b. After tracing through a series of diagrams, one will see that they're left with a torus with a hole removed. The four dashed arcs are all glued together and they become the event horizon.

2.4 3D Gravity Partition Function

In this section we'll provide a brief sketch of the three dimensional gravity partition function and then review a motivating example: the Hawking-Page phase transition (Hawking and Page, 1983). We'll go into more detail about the partition function in the context of AdS/CFT in section 4.2. A nice review of the three dimensional partition function in quantum gravity can be found in (Maloney and Witten, 2010).

We know that the general form of the partition function will be $Z = \int e^{iI}$, where I is the action for the theory of interest and the integral is over all solutions that satisfy the appropriate boundary conditions. Eventually we'll consider quantum field theory in curved spacetimes and the partition functions will need to sum over both field and spacetime variations. As a warmup we'll first recall the partition function in quantum field theory. In quantum field theory the partition function is given by $Z = \int \mathcal{D}\phi e^{iI[\phi]}$. It's not guaranteed that these integrals will converge so one normally performs a Wick rotation $\tau = it$ to obtain the Euclidean path integral

$Z = \int \mathcal{D}\phi e^{-I[\phi]}$, which converges.

We can connect this to statistical mechanics by thinking of this as a trace over states. For time-invariant systems the transition amplitude to go from field configuration ϕ_1 at time t_1 to field configuration ϕ_2 at time t_2 can be written as $\langle\phi_2|e^{-iH(t_2-t_1)}|\phi_1\rangle$. If we set $\phi_1 = \phi_2 = \phi$ and Wick rotate so that $i(t_2 - t_1) = \tau_2 - \tau_1 := \beta$, and then integrate over all possible field configurations, we obtain

$$Z = \mathrm{Tr}\exp(-\beta H), \tag{2.52}$$

where the path integral is now taken over all fields that are periodic in imaginary time with period β. In other words this is the partition function at inverse temperature β. The dominant contribution to the partition function is given by the minima of the action where $\delta I = 0$.

In curved spacetime the partition function needs to sum over variations in spacetime geometry in addition to fields. Consider a field ϕ propagating from an initial configuration ϕ_1 on a two-dimensional surface Σ_1 with metric g_1 to a field ϕ_2 on surface Σ_2 and metric g_2. The partition function is given by

$$Z = \int \mathcal{D}[g, \phi]e^{iI[\phi]} \tag{2.53}$$

where the path integral is taken over all fields ϕ and spacetime manifolds $\mathcal{M}$ so that the fields have the correct initial and final configurations and the spacetime manifolds have the correct initial and final time slicings (with appropriate induced metric.) We perform a Wick rotation to get a convergent path integral (but now with Euclidean contributions)

$$Z = \int \mathcal{D}[g, \phi]e^{-I[\phi]}. \tag{2.54}$$

The action is the usual Euclidean version of the Einstein-Hilbert action

$$I = \int_{\mathcal{M}} d^{n+1}x\sqrt{g}\frac{R - 2\Lambda}{16\pi G_N} \tag{2.55}$$

where G_N is Newton's constant, n are the number of spatial dimensions, g is the determinant of the metric tensor, R is the Ricci scalar, Λ is the cosmological constant and $\mathcal{M}$ is the spacetime manifold of interest. Note that the full action would also include a Gibbons-Hawking-York boundary term, but this won't be important in the Hawking-Page example, or in the rest of this paper so we didn't include it.

Hawking-Page phase transition

The Hawking-Page phase transition is more interesting in $3+1$ dimensions so we'll review it there. In this section we'll set $c = \hbar = G_N = 1$, but restore the AdS radius of curvature L_{AdS}. Before describing the phase transition in AdS lets recall some facts about $3+1$-dimensional Schwarzschild black holes. The temperature of a Schwarzchild black hole is given by the formula $T_H = \frac{1}{8\pi M}$ where M is the mass of the black hole. The specific heat is given by

$$\left.\frac{\partial E}{\partial T}\right|_V = \frac{\partial M}{\partial T} = -\frac{1}{8\pi M^2}. \tag{2.56}$$

The negative specific heat shows that Schwarzschild black holes are thermodynamically unstable.

However, things are more interesting in anti de-Sitter space, where AdS-black holes have a positive specific heat. Hawking and Page realized that although black holes can be in stable equilibrium with thermal equilibrium, below a certain critical temperature they are not the preferred state. AdS_4 is the maximally symmetric solution of Einstein's equations with constant negative curvature. Set $L_{AdS} = \ell$ for convenience of notation. We can think of AdS_4 as a submanifold in $\mathbb{R}^{2,3}$ given by the constraint

$$AdS_4 = \left\{(x_0, x_1, x_2, x_3, x_4) \in \mathbb{R}^{2,3} \mid -x_0^2 - x_4^2 + x_1^2 + x_2^2 + x_3^2 = -\ell^2\right\}. \tag{2.57}$$

As with AdS_3 there are many different coordinate systems that can be put on this manifold. Without going through all the details we'll only need the 'static' coordinates which are obtained via a series of substitutions. First set

$$x_0 = \ell \cosh\rho \cos\bar{t}, \quad x_4 = \ell \cosh\rho \sin\bar{t}, \quad x_i = \ell \sinh\rho\Omega_i \tag{2.58}$$

where $\bar{t} \in S^1$, $\rho > 0$, $i \in \{1,2,3\}$ and $\sum_i \Omega_i = 1$. The metric becomes

$$ds^2 = \ell^2\left(-\cosh^2\rho d\bar{t}^2 + d\rho^2 + \sinh\rho d\Omega_2^2\right) \tag{2.59}$$

where $d\Omega_2^2$ is the spherical metric on S^2. These coordinates have closed time-like curves which can be removed by lifting $\bar{t} \in S^1$ to $\bar{t} \in \mathbb{R}$. The static coordinates are then obtained by making the substitution $t = \ell\bar{t}$ and $r = \ell\sinh\rho$ so that the metric becomes

$$ds^2_{AdS_4} = -\left(1 + \frac{r^2}{\ell^2}\right)dt^2 + \left(1 + \frac{r^2}{\ell^2}\right)^{-1} dr^2 + r^2 d\Omega_2^2. \tag{2.60}$$

We also need the metric for AdS-Schwarzschild black holes which is given by

$$ds^2_{AdS-BH} = -\left(1 - \frac{2M}{r} + \frac{r^2}{\ell^2}\right) dt^2 + \left(1 - \frac{2M}{r} + \frac{r^2}{\ell^2}\right)^{-1} dr^2 + r^2 d\Omega_2^2. \tag{2.61}$$

Note that at small r this metric looks like the typical Schwarzschild metric in Minkowski space and at large r it looks like the metric for AdS_4. The AdS black hole has an event horizon at $r = r_+$ where r_+ is the largest root of $V(r) = \left(1 - \frac{2M}{r} + \frac{r^2}{\ell^2}\right)$. The temperature of the AdS black hole is given by

$$T_{BH} = \frac{\ell^2 + 3r_+^2}{4\pi\ell^2 r_+}. \tag{2.62}$$

This value can be obtained either by calculating the surface gravity in the Lorentzian solution, or analytically continuing via a Wick rotation $\tau = it$, and finding the period of τ so that there isn't a conical singularity at the horizon. The black hole temperature is the inverse of this period, which we'll call β_{BH}

$$\beta_{BH} = \frac{4\pi\ell^2 r_+}{\ell^2 + 3r_+^2}. \tag{2.63}$$

Note that the black hole temperature is not monotonically decreasing as a function of mass. There's a critical temperature $T_0 = \frac{\sqrt{3}}{2\ell\pi}$ below which AdS-Schwarzschild black holes can't exist.

We're now in a position to compare contributions to the partition function. Note that by contracting Einstein's field equations with the inverse metric tensor we see that $R = 4\Lambda$. After plugging this back into the Einstein-Hilbert action we see that the classical action reduces to volume integrals

$$I = \frac{\Lambda}{8\pi} \int d^4x \sqrt{-g}. \tag{2.64}$$

We can ignore the Gibbons-Hawking-York boundary terms because they'll cancel when we consider the difference $I_{BH} - I_{AdS}$. However, these volume integrals are each infinite so we need to regularize them somehow. The standard way is to put a cut-off surface in the radial direction, so that instead of integrating r all the way to infinity we only integrate to $r = K$. The contributions become

$$I_{AdS} = \frac{\Lambda}{8\pi} \int_0^{\beta_{AdS}} d\tau \int_0^K r^2 dr \int_{S^2} d\Omega^2 = \frac{\Lambda K^3}{6} \beta_{AdS} \tag{2.65}$$

for the AdS metric and

$$I_{BH} = \frac{\Lambda}{8\pi} \int_0^{\beta_{AdS}} d\tau \int_{r_+}^K r^2 dr \int_{S^2} d\Omega^2 = \frac{\Lambda(K^3 - r_+^3)}{6} \beta_{BH}. \tag{2.66}$$

β_{BH} is fixed by the formula 2.63 but β_{AdS} can take any value. We'll fix it by requiring the metrics to agree at the $r = K$ hypersurface, where

$$ds^2_{AdS_4} = \left(1 + \frac{K^2}{\ell^2}\right) d\tau^2 + K^2 d\Omega_2^2 \tag{2.67}$$

$$ds^2_{AdS-BH} = \left(1 - \frac{2M}{K} + \frac{K^2}{\ell^2}\right) d\tau^2 + K^2 d\Omega_2^2. \tag{2.68}$$

In particular, these will match if the time coordinates have the same period which requires

$$\beta_{AdS} = \beta_{BH} \frac{\sqrt{1 - \frac{2M}{K} + \frac{K^2}{\ell^2}}}{\sqrt{1 + \frac{K^2}{\ell^2}}}. \tag{2.69}$$

Now if we look at the difference $I := I_{BH} - I_{AdS}$ and take the limit as $K \to \infty$ we obtain

$$I = I_{BH} - I_{AdS} = \frac{\pi r_+ (\ell^2 - r_+^2)}{\ell^2 + 3r_+^2}. \tag{2.70}$$

In this form it's clear that when $r_+ < \ell$ the AdS_4 contribution is smaller so it will be the dominant contribution in the partition function. When $r_+ > \ell$ the black hole contribution dominates the partition function. When $r_+ = \ell$ there is a first order phase transition. The temperature when these phases exchange dominance is given by plugging $r_+ = \ell$ back into the formula for an AdS black hole's temperature (2.62) to obtain the critical temperature: $T_H = \frac{1}{\pi \ell}$.

Chapter 3

Spectral Theory Background

This section goes into more detail than is needed, but we figured it might be valuable for the physics community to summarize the key results from the relevant mathemetics literature. This section draws heavily from (McMullen, 1999) and (Borthwick, 2007) which organized results of (Sullivan, 1979), (Sullivan et al., 1987)) and (Bishop and Jones, 1997), into a form which makes the relationship between the Hausdorff dimension of the limit set, $\mathrm{H.dim}(\Lambda(\Gamma))$, and the first resonance of the Laplacian, $\lambda_0(\Gamma)$, particularly clear.

To set some notation, let $\mathbb{H}^{d+1}$ be the upper half-plane model of hyperbolic space with constant curvature -1 and let $S^d_\infty = \mathbb{R}^d \cup \{\infty\}$ denote its boundary. Note that all of the following results hold for any model of hyperbolic space (and in particular these quantities will be invariant under conformal maps).

3.1 Schottky groups

Not all Kleinian groups give us spacetimes that have a physical interpretion. The groups of interest live in a subgroup of $PSL(2,\mathbb{C})$ and they are called Schottky groups. Schottky groups are Kleinian groups such that each generator, call it A, maps a circle C_A to another disjoint circle C_a. Thinking of A as a Mobious transformation, A maps the interior of C_A to the exterior of C_a. We will use lower-case/upper-case letters to denote transformations and their inverses respectively. In our notation

$$A : C_A \mapsto C_a \tag{3.1}$$

$$a : C_a \mapsto C_A \tag{3.2}$$

$$\tag{3.3}$$

which should be clear from figure 2.2 and figure 2.3. Note that there are two types of Schottky groups: 'classical' and 'non-classical.' Classical Schottky groups have a finite number of generators and these generators pair honest-to-goodness circles to other circles. All of the circles need to be disjoint from each other. Non-classical Schottky groups can have infinitely many generators and the circle pairings can be to objects that are only homeomorphic to circles, such as pairing closed Jordan curves. For non-classical Schottky groups we'll still need the condition that all of the "circles" are disjoint from each other. Non-classical Schottky groups still give physically relevant spacetimes; they correspond to solutions that have spinning black holes in the bulk, or correspondingly, angular momentum on their boundary.

There's a trick from "Indra's Pearls" that makes it easy to find new Schottky groups given a starting Schottky group. They call it 'the power of conjugation' and itt comes from the following simple algebra. If $A \in PSL(2,\mathbb{C})$ maps the circle C_A to the circle C_a and G is an arbitrary element of $PSL(2,\mathbb{C})$ then GAG^{-1} maps the circle GC_A onto the circle GC_a. By GC_A we mean taking every point on the circle C_A and finding its image by G after fractional linear transformation. It's convenient to denote the circles GC_A and GC_a as C_{GA} and C_{Ga} respectively. We'll use this trick in section 5.2 to find new multiboundary wormholes given simple starting generators.

3.2 Schottky space

For each fixed genus $g \geq 2$, Schottky space $\mathcal{S}_g$ is the space of Schottky groups of genus g. Basically each point in Schottky space corresponds to a set of g elements of $PSL(2,\mathbb{C})$ that generate a Schottky group, up to Mobius transformations. This space has complex dimension $3g-3$ which follows from the fact that each Schottky group generates a genus g Riemann surface Σ_g. Σ_g can be decomposed into $2g-2$ pairs of pants, each of which is described by 3 complex degrees of freedom. But we glue all of the pairs of pants together so the number of degrees of freedom is $3g-3$. The universal covering space of $\mathcal{S}_g$ can be identified with the Teichmuller space of compact genus g Riemann surfaces. The notion of Schottky space provides a convenient way to keep track of the contributions one needs to sum over the three dimensional quantum gravity partition function. A more familiar parameterization of the space of Riemann surfaces of fixed genus g is given by the moduli space $\mathcal{M}_g$. Schottky space and moduli space are related in that they are both covered by Teicmuller space. Any of these parameterizations can be used in the partition function, but Schottky space makes it particularly easy to avoid double counting

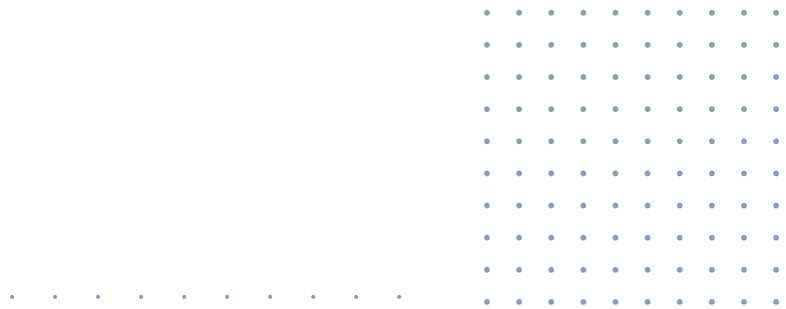

Figure 3.1: These figures show the "limit sets" that are obtained when lifting a point from $\mathbb{R}^2/\langle x \sim x+1\rangle$ and $\mathbb{R}^2/\langle x \sim x+1, y \sim y+1\rangle$, respectively, to their covering space $\mathbb{R}^2$. The limit sets create regularly spaced grids that have Hausdorff dimension equal to zero.

contributions.

3.3 Limit sets

For intuition, let's think about quotients of $\mathbb{R}^2$ for a moment. If one quotients $\mathbb{R}^2$ by the subgroup generated by translating x by coordinate distance one, $\langle x \sim x+1\rangle$, then the quotient $\mathbb{R}^2/\langle x \sim x+1\rangle$ is a cylinder with infinite length and circumference of length 1. If one quotients $\mathbb{R}^2$ by a subgroup generated by two translations, such as $\mathbb{R}^2/\langle x \sim x+1, y \sim y+1\rangle$, then one is left with a torus. Imagine drawing a point on both of these quotient spaces and then lifting to its covering space. In the cylinder case, the image of the point will be infinitely many points located at each of the integers along the real line. In the torus case, the image of the point will be an infinite 2d lattice with images at the tuples of integers $(a, b) \in \mathbb{Z} \times \mathbb{Z}$. See figure 3.1 for an illustration of this.

We started with these examples because they illustrate the concept of limit sets. Note that in the case of quotients of $\mathbb{R}^2$ the images of a point accumulate all over the plane and there's no sense of "stacking up." The fractal dimensions of limit sets in Euclidean space will be zero. However, when we repeat this lifting procedure in hyperbolic space, the points will accumulate near each other due to the negative curvature, creating infinitely nested fractals.

Let's now proceed to the definition of limit sets for Kleinian groups. Let $\Gamma \subseteq \mathrm{Isom}(\mathbb{H}^{d+1})$ be a Kleinian group. The *limit set* $\Lambda(\Gamma)$ is the subset of the boundary

sphere $S^d_\infty = \mathbb{R}^d \cup \{\infty\}$ defined for any $x \in \mathbb{H}^{d+1}$ by

$$\Lambda(\Gamma) = \overline{\Gamma x} \cap S^d_\infty. \tag{3.4}$$

This is equivalent to taking the union of the limits of orbits, $g^n x$ as $n \to \infty$ for all words $g \in \Gamma$, where g^n acts on x by fractional linear transformation when $d = 2$. This definition corresponds exactly to the intuition you should have from the $\mathbb{R}^2$ case. You start with a point in the quotient space and then lift it to $\mathbb{H}^d$ where every image corresponds to the image of x by some reduced word $g \in \Gamma$. See figures 3.2 and 3.3 for examples of limit sets. These particular limit sets correspond to complicated spacetimes that have three asymptotic regions connected by a wormhole, where each of the throats is spinning rapidly. They were chosen so that the quotients of their boundaries, $\hat{\mathbb{C}}/\Gamma$, are genus two Riemann surface with a $\mathbb{Z}_3$ symmetry. These surfaces arise when investigating $n = 3$ Rényi entropies. Moreover, these two in particular correspond to non-classical Schottky groups with large Hausdorff dimension (greater than one).

3.4 Hausdorff dimension

In an attempt to be self-contained, this subsection introduces *Hausdorff dimension*, denoted $\mathrm{H.\,dim}(E)$ for the set E. Hausdorff dimension generalizes the concept of dimension to sets with non-integer dimension, such as fractals. The general definition is a bit involved, but it simplifies when the set of interest is self-similar, such as the sets we care about in this paper.

Let X be a metric space. If $E \subset X$ and $d \in [0, \infty)$, the d-dimensional Hausdorff content of E is defined as

$$C^d_H(E) := \inf_\delta \left\{ \sum_i r_i^d = \delta \quad \text{s.t.} \quad \text{there is a cover of } E \text{ by balls with radii } r_i > 0 \right\}. \tag{3.5}$$

With this definition in mind, the *Hausdorff dimension* of E is defined to be

$$\mathrm{H.\,dim}(E) := \inf_d \left\{ d \geq 0 \quad \text{s.t.} \quad C^d_H(E) = 0 \right\}. \tag{3.6}$$

For a self-similar set, E, the definition simplifies. We can consider coverings with the same sized balls. For balls of diameter ϵ, let N_ϵ be the number of balls we need to cover E. The Hausdorff dimension corresponds to the d where N_ϵ is proportional to ϵ^{-d}, so that for self-similar E, we have

$$\mathrm{H.\,dim}(E) = -\lim_{\epsilon \to 0^+} \frac{\ln N_\epsilon}{\ln \epsilon} \tag{3.7}$$

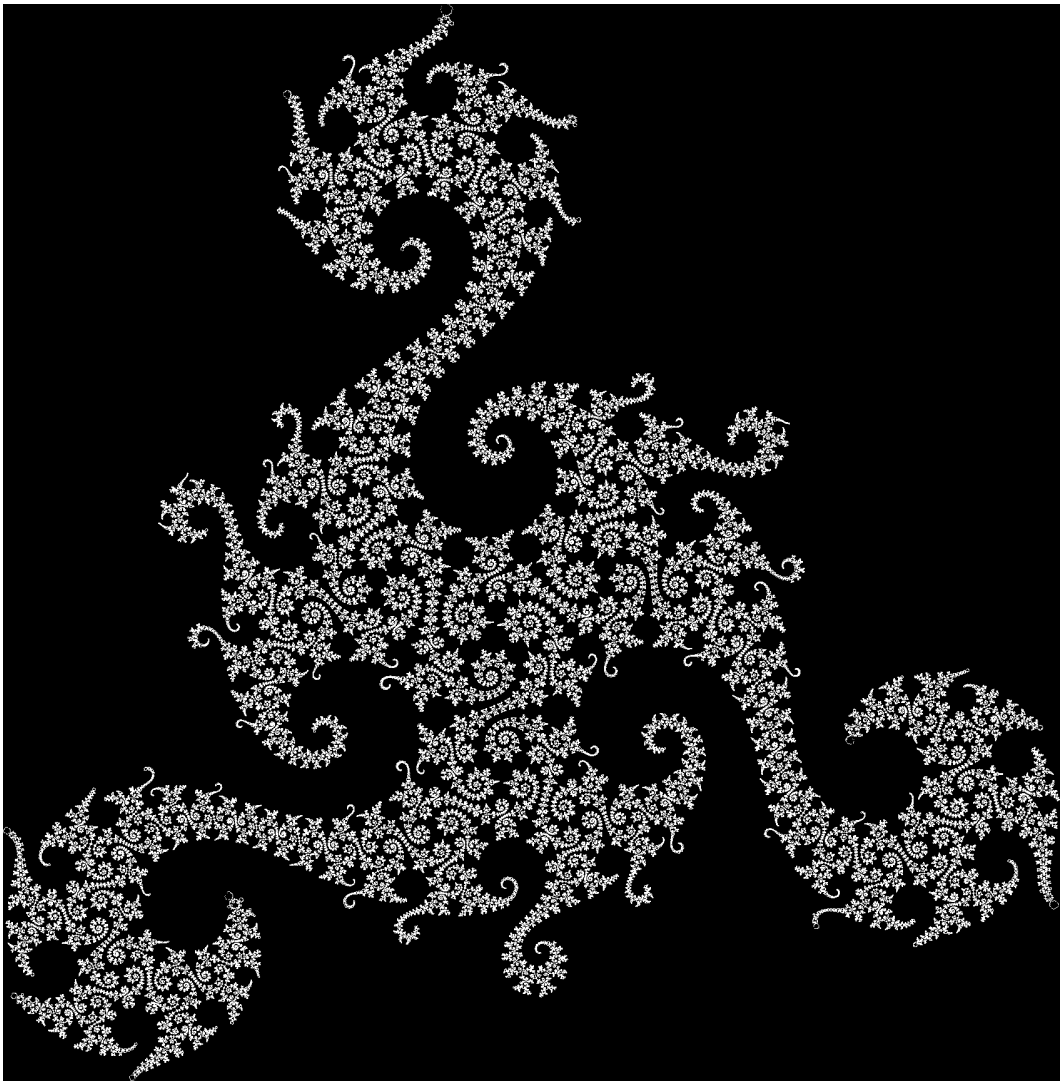

Figure 3.2: Example of a limit set generated by a Schottky group that corresponds to a Riemann surface with a $\mathbb{Z}_3$ symmetry. Using the parameterization 4.8 this has $q = .8 + .443i$. This Schottky group lives close to the boundary of Schottky space, where the limit sets are particularly rich. The Hausdorff dimension here is greater than 1.5. The $\mathbb{Z}_3$ symmetry results from this being a Riemann surface of interest for calculating $n = 3$ Rényi entropies.

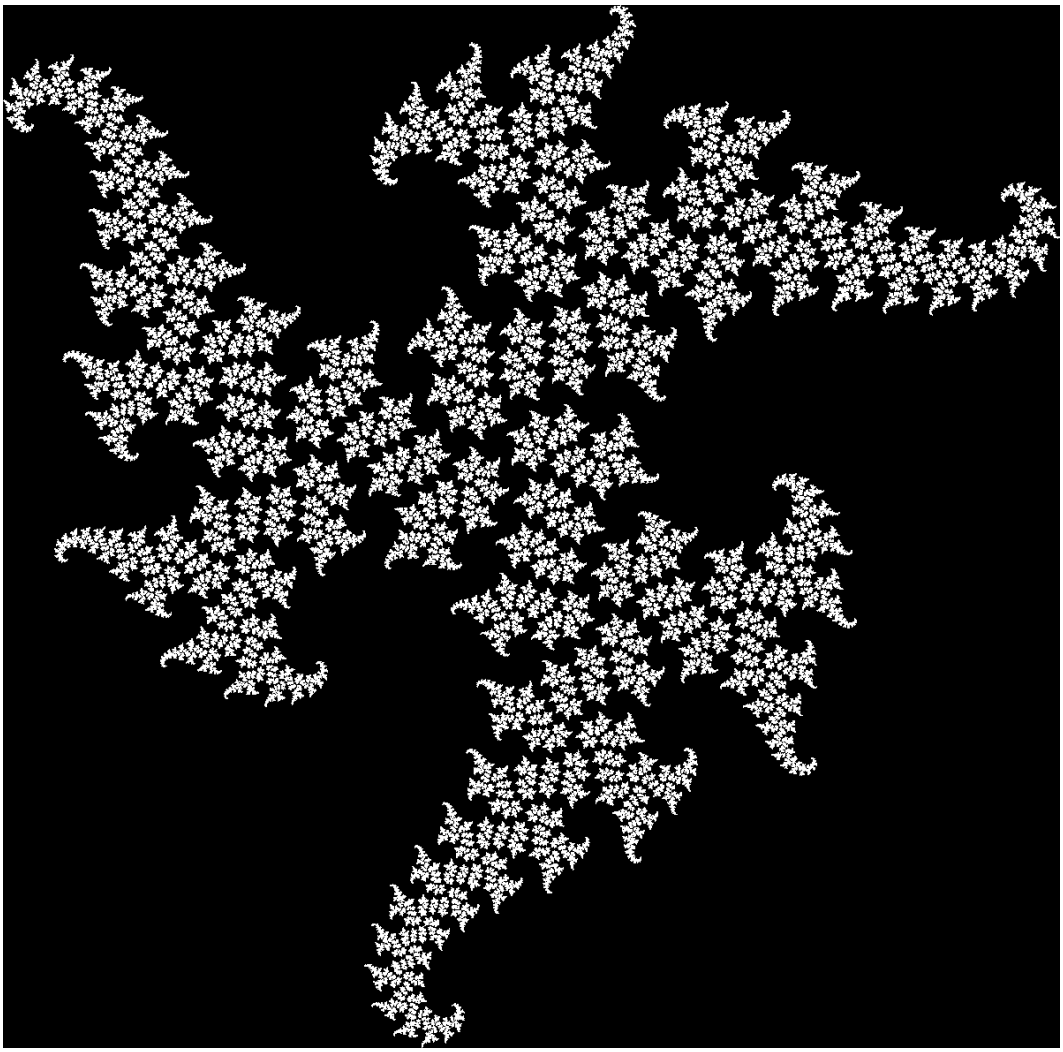

Figure 3.3: Example of a limit set generated by a Schottky group that corresponds to a Riemann surface with a $\mathbb{Z}_3$ symmetry. Using the parameterization 4.8 this has $q = .4 + .572i$. This Schottky group lives close to the boundary of Schottky space, where the limit sets are particularly rich. The Hausdorff dimension here is greater than 1.4.

Let's take the Cantor set as an example. At the n-th level we have 2^n line segments each of length 3^{-n}. Plugging in $\epsilon = 3^{-n}$ the formula reduces to

$$\text{H. dim}(E) = -\lim_{n\to\infty} \frac{\ln 2^n}{\ln 3^{-n}} = \frac{\ln 2}{\ln 3} \tag{3.8}$$

which recovers the well known value for the Hausdorff dimension of the Cantor set.

Calculating the Hausdorff dimension of limit sets of Schottky groups is usually difficult, but McMullen introduced a numerical algorithm in (McMullen, 1998) which we'll introduce later in this section and then work through examples in section 4.5.

3.5 First resonance of the Laplacian

Let us now go into more detail by recalling that the Laplacian on $\mathbb{H}^{d+1}$ is defined to be

$$\Delta := \frac{1}{\sqrt{|\det g|}} \partial_\mu \left(\sqrt{|\det g|} g^{\mu\nu} \partial_\nu \right) \tag{3.9}$$

which is invariant under isometries (for any metric)

$$\Delta \circ g = g \circ \Delta, \quad \text{for } g \in \text{Isom}(\mathbb{H}^{d+1}) \tag{3.10}$$

so it descends to a well defined Laplacian, Δ_Γ, on quotients $\mathbb{H}^{d+1}/\Gamma$.

It's worth mentioning that the mathematics and physics literatures use a different sign convention, which can oftentimes be confusing. Mathematicians add an extra minus sign in equation (3.9), which makes eigenvalues positive for positive eigenfunctions, $f \geq 0$. Physicists do not include this minus sign, which has the drawback that eigenvalues corresponding to L^2 functions are negative, but it has the advantage of keeping the expected sign in the wave equation. Also note that mathematicians sometimes call the first resonance the "bottom" of the Laplacian.

The *first resonance of the spectrum of the Laplacian* is defined to be

$$\lambda_0(\Gamma) := -\inf_{f\in C_0^\infty} \left\{ \frac{\int_{\mathbb{H}^{d+1}/\Gamma} |\nabla f|^2}{\int_{\mathbb{H}^{d+1}/\Gamma} |f|^2} \right\} \tag{3.11}$$

$$= -\sup_{f\in C_0^\infty} \left\{ \lambda \geq 0 \;\middle|\; \exists f > 0 \text{ on } \mathbb{H}^{d+1}/\Gamma \text{ with } \Delta f = \lambda f \right\} \tag{3.12}$$

$$= \inf_{f\in C_0^\infty} \left\{ \lambda \leq 0 \;\middle|\; \exists f > 0 \text{ on } \mathbb{H}^{d+1}/\Gamma \text{ with } \Delta f = \lambda f \right\} \leq 0 \tag{3.13}$$

where again, please note that this is the physics sign convention. Equation (3.12) is related to (3.13) by using $-\sup(x) = \inf(-x)$. The relationship between (3.11) and (3.12) was shown in (Cheng and Yau, 1975). This value, $\lambda_0(\Gamma)$, is important because it separates the spectrum between its L^2 and its 'positive parts.'

For a concrete example, let $d = 1$ so that $\Delta_{\mathbb{H}^2} := y^2(\partial_y^2 + \partial_x^2)$. The eigenfunctions are given by y^s with associated eigenvalue $s(s-1)$. The resonance is given by taking $s = \frac{1}{2}$ which implies $\lambda_0(\mathbb{H}^2) = -\frac{1}{4}$.

3.6 Critical exponent of the Poincaré series

The *Poincare Series* is defined for $x \in \mathbb{H}^{d+1} \cup \Omega(\Gamma)$ by

$$P_s(\Gamma, x) = \begin{cases} \sum_{g\in\Gamma} e^{-sd(x,gx)} & \text{if } x \in \mathbb{H}^{d+1} \\ \sum_{g\in\Gamma} |g'(x)|^s & \text{if } x \in \Omega(\Gamma) \end{cases} \tag{3.14}$$

where the derivative $g'(x)$ is measured in terms of the canonical spherical metric

$$\sigma = 2|dx|/(1 + |x|^2) \tag{3.15}$$

on S^d_∞. The Poincare series has a critical exponent for convergence given by

$$\delta(\Gamma) = \inf\left\{ s \geq 0 \;\middle|\; P_s(\Gamma, x) < \infty \right\} \tag{3.16}$$

which is independent of the x chosen. The Poincaré series is related to the method of images and to Greens functions.

3.7 Invariant conformal densities

A Γ-*invariant conformal density of dimension* α is a positive measure μ on S^d_∞ such that

$$\mu(gE) = \int_E |g'|^\alpha d\mu \tag{3.17}$$

for every Borel set E and $g \in \Gamma$. A density is normalized if $\mu(S^d_\infty) = 1$. This is just a fancy way of saying that it's a measure that will associate the same volume to sets related by isometries $g \in \Gamma$. Conformal densities aren't used directly in this paper but they're at the heart of how McMullen's numerical algorithm can efficiently estimate Hausdorff dimension. It's also sitting in the background because technically, when doing field theory in these quotient spaces, the measure used in the action will be a Γ-invariant conformal density. Hence, on a technical level they are

related to the coupling between bulk fields and boundary operators in AdS/CFT. The critical dimension of Γ is defined to be

$$\alpha(\Gamma) = \inf\left\{\alpha \geq 0 \;\middle|\; \exists \text{ a } \Gamma\text{-invariant density of dimension } \alpha\right\}. \tag{3.18}$$

Note that there is an equivalent definition using the functorial point of view: a conformal density of dimension α is a map

$$\mu : (\text{ conformal metrics } \rho(z)|dz| \text{ on } S^d_\infty) \to (\text{measures on } S^d_\infty) \tag{3.19}$$

such that

$$\frac{d\mu(\rho_1)}{d\mu(\rho_2)} = \left(\frac{\rho_1}{\rho_2}\right)^\alpha. \tag{3.20}$$

Conformal maps acts on densities in a natural way and (3.17) shows that $g_*(\mu) = \mu$ for any conformal map g.

3.8 Summary of key spectral theory results

With all of these definitions in place we're now ready to summarize the key results from the mathematics literature. For any nonelementary complete hyperbolic manifold $M = \mathbb{H}^{d+1}/\Gamma$ the following are equivalent

1. The Hausdorff dimension, H. $\dim(\Lambda(\Gamma))$, of the limit set of Γ
2. The critical exponent, $\delta(\Gamma)$, of the Poincaré series,
3. The minimum dimension, $\alpha(\Gamma)$, of a Γ-invariant conformal density
4. The asympotic length spectrum of the prime geodesics: if $\pi(\ell) := |\{\gamma \in \mathcal{P} : \ell_\gamma < \ell\}|$ is the prime geodesic counting function, we have the prime geodesic theorem
$$\pi(\ell) \sim \frac{e^{\delta\ell}}{\delta\ell} \tag{3.21}$$
(This can be improved to $\pi(\ell) = \mathrm{li}(e^{\delta\ell}) + O(e^{(\delta-\epsilon)\ell})$ for some positive ϵ.)
5. The first zero/pole in various definitions of Selberg zeta functions

Moreover, the first resonance of the Laplacian is given in terms of the above quanities via the relationship

$$\lambda_0(\Gamma) = \begin{cases} -\frac{d^2}{4} & \text{if } \delta(\Gamma) \leq d/2, \\ \delta(\Gamma)(\delta(\Gamma) - d) & \text{if } \delta(\Gamma) \geq d/2 \end{cases} \tag{3.22}$$

where δ is any of the quantities above but usually interpreted as the Hausdorff dimension of the limit set. This relationship is central in this paper because it shows that we can gain information about the spectrum of the Laplacian from geometrical quantities that are oftentimes easier to compute.

3.9 McMullen's Algorithm

Given a Euclidean spacetime obtained as a quotient $\mathbb{H}^{d+1}/\Gamma$, calculating the spectrum of the Laplacian is often difficult to do directly. Coming at this question from a different direction, McMullen introduced a numerical algorithm that efficiently calculates the associated Hausdorff dimension, which indirectly provides information about the spectrum via the results in appendix 3.8. This subsection introduces McMullen's algorithm (McMullen, 1998).

The algorithm works by iteratively approximating a Γ-invariant conformal density. Appendix 3.7 goes into more detail, but recall that a Γ-invariant conformal density with dimension α is a probability measure μ on $\partial\mathbb{H}^{d+1}$ such that

$$\mu(\gamma E) = \int_E |\gamma'|^\alpha d\mu \tag{3.23}$$

for any Borel set E and $\gamma \in \Gamma$. Think of this as a probability measure that gives consistent values for two subsets of the limit set that only differ by an isometry γ. There's a result due to Sullivan showing that the limit set of a geometrically finite Kleinian group has a *unique* nonatomic Γ-invariant density μ of total mass one. This is important because the dimension of this unique density is equal to the Hausdorff dimension of the limit set by the theorem in appendix 3.8.

McMullen's algorithm works by starting with a coarse-grained superset that covers the limit set and iteratively looking at finer and finer approximations. We'll restrict ourselves to a description of the algorithm for classical Schottky groups where the algorithm is simpler to describe. Moreover, we'll consider groups that don't have any parabolic elements (these are Mobius transformations that only have one fixed point and therefore correspond to cusps in the quotient space).

The algorithm starts by finding a Markov partition $\langle (D_i, f_i) \rangle$ for the group Γ. A Markov partition is a set of connected compact balls $D_i \subset S^n$ and maps $\gamma_i \in \Gamma$ defined on P_i such that:

1. $\gamma_i(D_i) \supset \cup_{i \mapsto j} D_j$, where the relation $i \mapsto j$ means $\mu(\gamma_i(D_i) \cap D_j) > 0$;

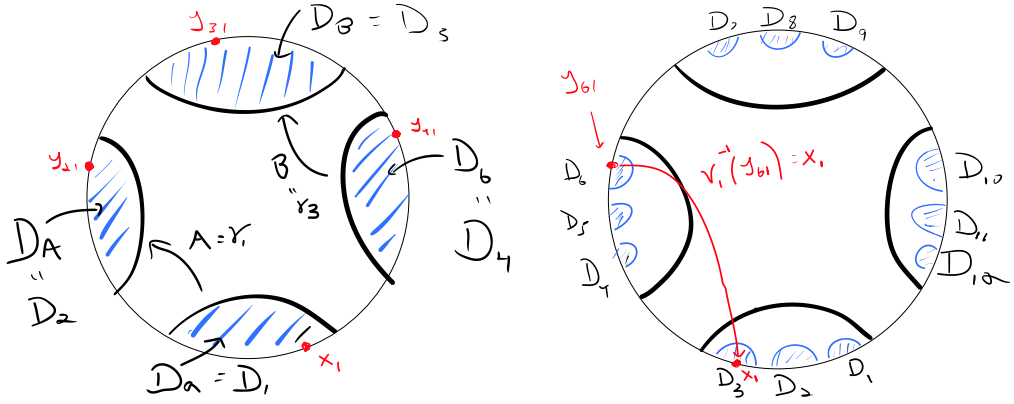

depth one depth two

Figure 3.4: This figure illustrates a Markov partition for a Schottky group with two generators corresponding to a pair of pants quotient. The Markov partition at the first level has four disks, $D_1 = D_a, D_2 = D_A, D_3 = D_B$ and $D_4 = D_b$, with corresponding maps A, a, B, b. The points y_{21}, y_{31} and y_{41} correspond to the three images of the sample point x_1. When checking the definition of a Markov partition remember that Schottky groups are defined so that the interiors of disks are mapped to the exterior of their target disk. For example, in this picture A maps the interior of D_a to the exterior of D_A.

2. γ_i is a homeomorphism of a neighborhood of $D_i \cap \gamma_i^{-1}(D_j)$, when $i \mapsto j$;
3. $\mu(D_i) > 0$;
4. $\mu(D_i \cap D_j) = 0$ if $i \neq j$; and
5. $\mu(\gamma(D_i)) = \mu(\cup_{i \mapsto j} P_j) = \sum_{i \mapsto j} \mu(P_j)$,

where μ is the canonical measure (spherical measure if you're considering $\partial \mathbb{H}^d$ to be a sphere; Euclidean measure if you're considering it to be the complex plane with point at infinity).

All this notation can be intimidating at first, but the ideas are simply illustrated with a picture, such as figure 3.4. This example illustrates a Schottky group $\langle A, B \rangle$ with two generators. The generator A maps the interior of D_a to the exterior of D_A and the generator B maps the interior of D_b to the exterior of D_B. The quotient under these maps gives a pair of pants. The most natural initial Markov partition is given by $\langle (D_a, A), (D_A, a), (D_b, B), (D_B, b) \rangle = \langle (D_i, \gamma_i) \rangle_{i=1,4}$. It should be clear that this satisfies the definition of a Markov partition after chasing through a few pictures.

There are two more ingredients that we need to describe before introducing the algorithm. In addition to starting with a Markov partition one also needs sample

points $x_i \in D_i$. These will be important because we're going to keep track of how the distance between points changes as we apply the γ_i. We also need to introduce a *refinement* of the Markov partition $\langle (D_i, \gamma_i) \rangle$ which is defined to be

$$\langle (D_{ij}, \gamma_i) \rangle \tag{3.24}$$

for $i \mapsto j$ and where

$$D_{ij} = \gamma_i^{-1}(D_j) \cap D_i. \tag{3.25}$$

In the language of Schottky groups this translates to taking each disk and breaking it into disks that are one depth smaller (where by depth we mean disks corresponding to words with one extra letter). Note that we don't change the maps γ_i at each level, they're the same as what we started with. Figure 3.4 shows how this works for a pair of pants geometry. The four initial disks get split into 12 disks. At level $n+1$ there will be $4 \cdot 3^n$ disks. For Schottky groups with k generators at the $n+1$-th level there will be $2k \cdot (2k-1)^n$ disks.

With these definitions in hand we can write down McMullen's algorithm to compute the dimension δ of the density μ:

1. For each $i \mapsto j$, solve for $y_{ij} \in D_i$ such that $\gamma_i(y_{ij}) = x_j$;
2. Compute the transition matrix

$$T_{ij} = \begin{cases} \left|\gamma_i'(y_{ij})\right|^{-1} & \text{if } i \mapsto j, \\ 0 & \text{otherwise} L \end{cases} \tag{3.26}$$

3. solve for $\alpha \geq 0$ such that the spectral radius of the matrix $[T_{ij}^{\alpha}]$ is equal to one. Note that here T_{ij}^{α} is entry-wise exponentiation rather than matrix exponentiation!
4. Output α as an approximation to the Hausdorff dimension;
5. Replace the Markov partition with its refinement and iterate steps 1-5.

McMullen was able to show that this algorithm converges to the Hausdorff dimension. Specifically he showed that at most $O(N)$ refinements are needed to compute the Hausdorff dimension to N digits of accuracy.

Note that McMullen's algorithm isn't guaranteed to work when Γ has parabolic elements. The necessary and sufficient condition is that the invariant density μ can't

have any ‘atoms’ at the cusp points (basically non-zero measure at the cusps.) If it can be argued that the Γ-invariant density doesn't have any atoms then McMullen's algorithm will apply, albeit with ‘worst-case’ convergence properties. McMullen used these ideas to approximate the Hausdorff dimension of the Apollonian gasket in (McMullen, 1998). This subtlety about parabolic elements is important when investigating the $n \to \infty$ limit of Rényi entropies because the quotient surface has cusps.

Chapter 4

Phase Transitions in AdS_3

This chapter is heavily based on [DMMM]. In this chapter we show that for three dimensional gravity with higher genus boundary conditions, if the theory possesses a sufficiently light scalar, there is a second order phase transition where the scalar field condenses. This three dimensional version of the holographic superconducting phase transition occurs even though the pure gravity solutions are locally AdS_3. This is in addition to the first order Hawking-Page-like phase transitions between different locally AdS_3 handlebodies. This implies that the Rényi entropies of holographic CFTs will undergo phase transitions as the Rényi parameter is varied, as long as the theory possesses a scalar operator which is lighter than a certain critical dimension. We show that this critical dimension has an elegant mathematical interpretation as the Hausdorff dimension of the limit set of a quotient group of AdS_3, and use this to compute it, analytically near the boundary of moduli space and numerically in the interior of moduli space. We compare this to a CFT computation generalizing recent work of Belin, Keller and Zadeh, bounding the critical dimension using higher genus conformal blocks, and find a surprisingly good match.

4.1 Outline

Three dimensional gravity has proven a remarkably rich testing ground for our ideas about classical and quantum gravity. Even though Einstein gravity possesses no local degrees of freedom, three dimensional theories of gravity nevertheless have many of the rich features of their higher dimensional cousins, including holography (J. M. Maldacena, 1999) and black hole solutions (Banados, Teitelboim, and Zanelli, 1992) whose Bekenstein-Hawking entropy can be computed microscopically (Strominger, 1998). Theories of gravity in AdS_3 are dual to two dimensional conformal field

theories, allowing one to use CFT methods to gain insight into classical and quantum gravity in AdS. In this chapter we will use CFT methods to motivate the existence of a new class of phase transitions in three dimensional gravity. We will then verify their existence directly in classical AdS gravity, and explore their features.

Our central result is simple, and is easiest to state for AdS_3 gravity in Euclidean signature. Such theories can be defined with a variety of boundary conditions: one can take the boundary of (Euclidean) space-time to be any smooth, two dimensional Riemann surface $\mathcal{B}$. With appropriate boundary conditions (Brown and Henneaux, 1986) the theory will depend only on the conformal structure of $\mathcal{B}$, so can be studied as a function of the conformal structure moduli of $\mathcal{B}$. The bulk gravity path integral with these boundary conditions is, via AdS/CFT, equal to partition function of the dual CFT on the surface $\mathcal{B}$. We will be interested in gravity theories in the semi-classical limit, where this bulk path integral is dominated by the classical geometry which minimizes the (appropriately regularized) gravitational action. For example, when the boundary is a sphere the dominant contribution comes from Euclidean AdS_3, i.e. hyperbolic space $\mathbb{H}_3$, which is the unique constant negative curvature metric on the solid ball which "fills in" the boundary sphere. At higher genus, however, this path integral can have many saddle points, each of which correspond to a gravitational solution whose boundary is the surface $\mathcal{B}$. For example, when the boundary is a torus the bulk saddles are constant negative curvature metrics on a solid donut which fills in the boundary torus. There are many such saddles, which are distinguished by which cycle in the boundary torus is contractible in the bulk (see e.g. (J. M. Maldacena and Strominger, 1998; Dijkgraaf et al., 2000; Maloney and Witten, 2010)). For example, the geometry for which the Euclidean time coordinate is contractible is the (Euclidean) BTZ black hole, while the geometry where the angular coordinate is contractible is interpreted as the "thermal AdS" geometry used to compute finite temperature observables in a fixed AdS background. As one varies the moduli of the torus, these two saddles will interchange dominance in the bulk gravity path integral – this is the three dimensional version of the Hawking-Page phase transition (Hawking and Page, 1983) describing black hole formation in AdS.

We are interested in the case where the boundary $\mathcal{B}$ has genus $g \geq 2$. Just as in the torus case, there are many different bulk solutions which give saddle point contributions to the partition function, and which can be characterized by a choice of cycles of the boundary surface $\mathcal{B}$ which become contractible in the bulk. The

simplest of these solutions are handlebodies, where the bulk solution is the constant negative curvature metric on a solid genus g surface which fills in the boundary $\mathcal{B}$. There will be phase transitions where these geometries interchange dominance: these are the higher genus versions of the Hawking-Page phase transition. As in the torus case, these handlebodies can be regarded as the Euclidean continuation of AdS_3 black holes; they are analytic continuations not of the BTZ black hole, but instead of multi-boundary black holes in AdS (D. R. Brill, 1996; Aminneborg, Bengtsson, D. Brill, et al., 1998), as described in (Krasnov, 2000). In the holographic context, these handlebodies describe contributions to the higher genus partition functions of holographic CFTs, which can be used to compute entanglement Rényi entropies (Headrick, 2010), to constrain OPE coefficients (J. Cardy, Maloney, and Maxfield, 2017; Cho, Collier, and Yin, 2017a; Keller, Mathys, and Zadeh, 2017), or as models of multi-party holographic entanglement (Balasubramanian et al., 2014).[1]

The bulk solutions described above are all locally $\mathbb{H}^3$, so can be written as quotients of hyperbolic space of the form $\mathbb{H}^3/\Gamma$, where Γ is a discrete subgroup of the isometry group of $\mathbb{H}^3$. Indeed, Einstein gravity in three dimensions has no local degrees of freedom, so any solution of pure Einstein gravity must be locally $\mathbb{H}^3$. We are interested in more complicated theories of gravity, however, which have additional degrees of freedom. In this paper we will consider theories where we have an additional scalar field ϕ of mass m^2. This means that the dual CFT has an operator $\mathcal{O}$ of dimension Δ, with $m^2 = \Delta(\Delta - 2)$.[2] All of the solutions described above have $\phi = 0$, and are dual to CFT configurations on the Riemann surface with $\langle \mathcal{O} \rangle_{\mathcal{B}} = 0$.

Our central result is the following: in some regions of moduli space, and for sufficiently light scalar fields, the handlebody solutions described above are unstable. This is because the kinetic operator $(\nabla^2 - m^2)$ will have a negative eigenvalue. Thus the solution with least action will not be a quotient of AdS, but rather a non-Einstein solutions with $\phi \neq 0$. In the dual CFT, the expectation value of the scalar operator $\langle \mathcal{O} \rangle_{\mathcal{B}} \neq 0$ will be non-zero. This means that as the moduli are varied there will be phase transitions where these scalar fields condense. Although the general structure of these $\phi \neq 0$ solutions is quite complicated – we expect the construction of these

[1] There are other "non-handlebody" solutions as well (Yin, 2008b; Yin, 2008a), which will not concern us in this paper. Near boundaries of moduli space – i.e. where cycles in the surface $\mathcal{B}$ become small – the handlebody solutions will always dominate (Yin, 2008b; Yin, 2008a). Moreover, one can compute numerically the action of the non-handlebodies in the interior of moduli space, and – at least in the cases which have been studied – they are always subdominant compared to handlebodies (Maxfield, Ross, and Way, 2016). We will therefore focus only on handlebodies in this paper.

[2] We are working in units where the AdS radius is $L_{AdS} = 1$.

solutions to be a difficult numerical problem – we are able to prove rigorously the existence of the instability.

We will see that these phase transitions have several important features, including:

Instabilities occur only when the dual operator is sufficiently light. In order for a given handlebody to be unstable, Δ must be lighter than a certain critical value Δ_c which we will compute. The value of Δ_c will depend on the conformal structure moduli as well as on the choice of handlebody. Whenever the dimension $\Delta < 2$ (i.e. the bulk scalar ϕ has $m^2 < 0$) there is some region of moduli space where a given handlebody will be unstable.[3]

Instabilities occur only in the interior of moduli space. At the boundary of moduli space one of the handlebody phases will always dominate, and will be stable against the condensation of any scalar field. In other words, $\Delta_c \to 0$ for the dominant handlebody as we approach the edge of moduli space. As one moves into the interior of moduli space Δ_c increases so the handlebody becomes more unstable to condensation of the scalar, until the first-order Hawking-Page transition is reached, and a topologically distinct handlebody becomes dominant.

For example, if $\mathcal{B}$ is a genus $g = n - 1$ surface constructed as an n-fold cover of the sphere branched over 4 points, parameterized by their cross-ratio x, the handlebody which dominates when $x \to 0$ will become more unstable as x is increased. For this geometry, Δ_c is a monotonically increasing function of x.[4]

Handlebodies become more unstable as the genus increases. Instabilities only occur when $\mathcal{B}$ has genus $g \geq 2$. If $\mathcal{B}$ is the genus $g = n - 1$ surface constructed as an n-fold cover of the sphere branched over 4 points, then if we hold the cross-ratio x of the four points fixed, the corresponding handlebodies will become more unstable as n is increased. In other words, Δ_c is a monotonically increasing function of n. As we take $n \to \infty$ with fixed x, Δ_c approaches a finite value which depends on x but is always greater than $\frac{1}{2}$.

[3] If we require the handlebody to be invariant under a $\mathbb{Z}_2$ time reflection symmetry, so that it can be Wick rotated to a real Lorenztian solution, then this condition becomes $\Delta < 1$; this would require that the scalar field have $-1 < m^2 < 0$ and be quantized with alternate boundary conditions.

[4] It is important here that we are referring to the handlebody which dominates at small x. For the handlebody which dominates as $x \to 1$, Δ_c will be a monotonically decreasing function of x between 0 and 1.

This new phase transition in three dimensional gravity is quite similar to the holographic superconducting phase transition in higher dimensions (Gubser, 2008; Hartnoll, Herzog, and Horowitz, 2008). There is, however, one crucial difference, which is that in the present case the solutions which become unstable are locally AdS_3, and have no external potentials (aside from metric moduli) turned on. Our instability occurs because of *global* properties of the handlebody, not due to any local properties of the metric.[5] Although at first sight surprising, our results are a three dimensional version of a famous fact in two dimensions: the spectrum of the hyperbolic Laplacian on a Riemann surface depends not just on the local structure of the metric (which is always hyperbolic) but also on the moduli of the Riemann surface. Indeed, this spectrum is the central object of interest in the study of arithmetic and quantum chaos, and many of our results are borrowed from this literature.

Our results have important implications for entanglement entropies in two dimensional conformal field theories. For any state in a two dimensional CFT, one can consider the reduced density matrix associated to a particular spatial region. The Rényi entropies $\frac{1}{1-n}\log \mathrm{Tr}\,\rho^n$ can then be used to characterize the spatial entanglement structure of this state. When the spatial region is collection of intervals, the Rényi entropy is – via the replica trick – equal to the partition function of a CFT on a higher genus Riemann surface whose genus depends on n (see e.g. (Calabrese and J. L. Cardy, 2004)). For example, the Rényi entropy for a pair of intervals in the vacuum state is equal to the partition function on a genus $g = n-1$ Riemann surface. The entanglement entropy is then computed by considering these Rényi entropies as an analytic function of n, and continuing to $n \to 1$. In this procedure one assumes that the entropies are analytic functions of n. We have seen, however, that in holographic CFTs the Rényi entropies can undergo a phase transition as n is varied, at some finite value of $n > 1$. Thus the replica method for computing entanglement (von Neumann) entropies must be treated with care.[6] For example, if we consider the Rényi entropies for a pair of intervals, two handlebodies will interchange dominance precisely at cross-ratio $x = 1/2$ (Headrick, 2010); this is also exactly where

[5] A rather similar phenomenon was observed in higher dimensions in (Belin and Maloney, 2016), where hyperbolic black holes in AdS_4 were observed to undergo similar phase transitions even though the solutions were locally AdS. As in the present case, the instability only arose because the hyperbolic black hole solutions differed globally from AdS_4. Thus modes of the scalar field which are not normally present (since they are non-normalizable in global AdS) suddenly become normalizable and lead to a genuine instability of the locally AdS solution.

[6] Similar phenomena were observed for spherical entangling surfaces in higher dimensional holographic CFTs in (Belin, Maloney, and Matsuura, 2013; Belin, Hung, et al., 2015), and for the three dimensional $O(N)$ model in (Metlitski, Fuertes, and Sachdev, 2009).

the Ryu-Takayanagi formula for entanglement entropy (Ryu and Takayanagi, 2006) will undergo a phase transition. Our results imply, however, that if the CFT has a sufficiently light operator then the Rényi entropies will undergo a phase transition at cross-ratio $x < 1/2$. For example, the $n = 2$ Rényi entropy will undergo a phase transition if the theory has an operator with dimension $\Delta < \Delta_c = 0.189124\cdots$. We note that this instability occurs for values of n which are strictly larger than one – we do not expect a non-analyticity in a neighborhood of $n = 1$. It would be interesting to revisit the arguments of (Faulkner, 2013; Hartman, 2013; Headrick, 2010; Lewkowycz and J. Maldacena, 2013) in this context.

Our results also make clear a sense in which higher genus CFT partition functions differ qualitatively from those on the sphere or torus. The torus partition function, for example, was shown by Hartman, Keller and Stoica to take a universal form at large central charge, provided one assumes that the spectrum of light states (i.e. those with dimension less than the central charge) does not grow too quickly (Hartman, Keller, and Stoica, 2014). This universal form is precisely that of a dual three dimensional theory of gravity which has a Hawking-Page transition between a thermal state and a BTZ black hole, and the sparseness condition is obeyed by any bulk local quantum field theory and even by string theories with string scale $\ell_{\text{string}} \lesssim \ell_{\text{AdS}}$. At higher genus, however, we see that additional phase transitions are generic, and occur even for duals of local quantum field theories in the bulk. Thus at higher genus there is no analogous “universal partition function" at large central charge.

While our discussion will be entirely in the context of three-dimensional gravity, similar phenomena will also occur in higher dimensions. The most direct analogue is with locally AdS spacetimes, to which almost everything generalizes straightforwardly. In particular, the critical dimension Δ_c for an instability can be shown to be equal to the Hausdorff dimension of an appropriate limit set, just as we will see below for the three dimensional case.[7] This is slightly less natural than in three dimensions, because solutions to Einstein’s equations need not be locally hyperbolic, so it is not clear when such geometries would dominate the path integral. More generally, the same mechanism of instability can apply, with global properties of the solution moving the critical mass above the naïve Breitenlohner-Freedman bound. Heuristically, scalars of negative mass squared can be stable because the reduction in action from the mass term in a finite region is compensated for by the positive contribution to the action from the gradient, required to match with the boundary

[7] Indeed, one can directly reinterpret the results of (Belin and Maloney, 2016) in this context.

conditions at infinity; it is important here that the volume of a region does not grow faster with size than its perimeter in negatively curved spaces. Without altering the local curvatures, nontrivial topology can upset this mechanism for stability by reducing the size of a region's boundary, and hence the gradient contribution to the action, for a given volume.

The discussion of the paper will be phrased in terms of the Euclidean solutions, but the results have interesting implications in Lorentzian signature. The relevant CFT states are defined by a Euclidean path integral on a Riemann surface with one or more boundaries, generalizing the familiar examples of the path integral on the disc preparing the vacuum state, and on the cylinder preparing the thermofield double state on two entangled copies of the CFT Hilbert space. This defines the state at $t = 0$, which can be evolved in Lorentzian time.

To find the semiclassical bulk dual of these states, we must first find the Euclidean solution that dominates the path integral on the 'Schottky double', the closed Riemann surface formed by gluing the surface to its mirror image along each of the boundaries, by construction producing a $\mathbb{Z}_2$-symmetric surface. The dominant solution is expected to respect this boundary time-reflection symmetry, so the bulk surface Σ fixed by the reflection acts as an initial data surface for Lorentzian evolution, and the quantum state of the bulk fields is the Hartle-Hawking wavefunction on Σ. For one possible solution, the $t = 0$ slice Σ is conformal to the original Riemann surface, describing a single-exterior black hole with topology hidden behind a horizon for a single-boundary case, or a multi-boundary black hole with an exterior region for each boundary, all joined by a non-traversable wormhole. Even in pure gravity, there are several phases of the dominant bulk solution, so depending on the moduli the bulk state can also be disconnected copies of pure AdS (but with fields in a state different from the vacuum), or something else. For a more detailed review of these states, see (Maxfield, 2015; Maxfield, Ross, and Way, 2016; Skenderis and Rees, 2011; Maloney, 2015).

Now, if there is a sufficiently relevant scalar operator in the CFT, there is an additional second-order phase transition to a dominant bulk solution with a nonzero classical value for the dual scalar field. This means that the initial data on Σ includes some scalar field configuration, which will evolve in time. The fact that these states are not stationary will then be visible even for a classical observer outside any horizon. When the phase includes a black hole, the scalar outside the horizon will rapidly decay away, falling into the black hole. A more interesting time evolution occurs

when the dual state includes copies of pure AdS, which may now include some scalar configuration. When the amplitude is small, as will be the case close to the transition, linearised evolution will suffice, with the field bouncing around periodically, but eventually nonlinearities will likely become important, with resonances between different modes. Perhaps the most likely evolution thereafter is a turbulent cascade to excite higher and higher frequency modes, with the solution nonetheless being regular for all time, as evidenced by numerical studies of a massless scalar interacting only gravitationally (Bizoń and Jałmużna, 2013). This is different from the situation in higher dimensions, in which a black hole forms after finite time; this cannot occur in three dimensions, because there is a finite energy threshold between the vacuum and the lightest black hole.

In 4.2 we will review briefly a few salient features of three dimensional gravity, as well as the necessary aspects of CFT on Riemann surfaces. In 4.3 we will give a CFT argument for the existence of an instability, inspired by recent results of Belin, Keller and Zadeh (Belin, Keller, and Zadeh, 2017). The main idea is that a free bulk scalar field is dual to a generalized free field in the boundary CFT, and we can compute the contribution of such a field to the higher genus CFT partition function. Using the asymptotic value of the OPE coefficients of multi-trace operators built from a generalized free field, along with higher genus conformal blocks in the appropriate regime, one can show that these contributions diverge when the field is sufficiently light, which signals the phase transition. This argument allows us to bound the critical dimension of the scalar field; for example, for the genus two handlebody relevant for the computation of the third Rényi entropy of two intervals at cross-ratio $x = \frac{1}{2}$, we find $\Delta_c \geq 0.189121\cdots$.

In 4.4 we will turn to the bulk instability. We will first review how the zero mode of the instability relates to various notions from the spectral theory of the Laplacian on a general bulk geometry $\mathcal{M}$. We then specialize to the main case of interest, for which $\mathcal{M}$ is a quotient of hyperbolic space $\mathbb{H}^3$ such as a handlebody, and find that the critical dimension has a rather beautiful mathematical interpretation. The quotient is by a group of Möbius maps, which has a limit set, a subset of the Riemann sphere (the boundary of Euclidean AdS_3). This limit set has a finite Hausdorff dimension $\delta > 0$, which is sometimes referred to as the fractal dimension of the limit set. This Hausdorff dimension is precisely equal to the critical dimension of the scalar field, $\Delta_c = \delta$. In other words, a scalar is unstable if and only if its dimension is less than

the Hausdorff dimension of the limit set. The calculation in 4.3 can therefore be regarded as a CFT estimate of this Hausdorff dimension, which provides explicit lower bounds on δ.

In 4.5 we turn to the explicit computation of the critical dimension, using an algorithm of McMullen for computing the Hausdorff dimension. We will use the algorithm to compute the critical dimension analytically, finding the asymptotic behaviour of δ as the boundary of moduli space is approached, for the handlebody which dominates the partition function. We also describe what happens to the instability at large genus. We will also use the algorithm to efficiently compute the critical dimension numerically. For example, for the genus two surface described above, the Hausdorff dimension is $\delta = 0.189124\cdots$, close to our CFT bound. We use the numerical data to provide plots of this critical dimension as a function of moduli, and as a function of genus.

4.2 Review of higher genus partition functions in 3D gravity and 2D CFT

In this section we will review the description of higher genus Riemann surfaces, and the construction of solutions to three-dimensional gravity with such boundaries, which can be interpreted as saddle points for the higher genus partition function of a holographic CFT. In particular, we describe a class of symmetric surfaces that we will use as examples. We will also review the interpretation of these partition functions in terms of Rényi entropies.

Moduli spaces and handlebodies

We are interested in studying holographic two-dimensional conformal field theories, dual to three dimensional AdS gravity in Euclidean signature, in particular on a Riemann surface $\mathcal{B}$ of genus $g \geq 2$. The partition function of the theory on such a surface, denoted $Z_g(\tau)$, will depend on the conformal structure of the surface $\mathcal{B}$.

Here τ is a collection of $3g-3$ complex coordinates which parameterize the moduli space $\mathcal{M}_g$ of conformal structures on $\mathcal{B}$.

At genus one, τ can be identified with the usual torus modulus. At higher genus there are various different coordinates which can be used to describe the moduli τ, some of which we will now describe.[8]

[8] Because of the conformal anomaly the partition function will in addition depend on a choice of metric within a given conformal class. Thus Z_g should not – strictly speaking – be regarded

For many purposes in CFT and gravity, the most convenient way to realize a Riemann surface $\mathcal{B}$ is as a quotient of the Riemann sphere $\mathbb{C}^*$ by a *Schottky group* Γ, $\mathcal{B} = \Omega(\Gamma)/\Gamma$. Here Γ is a discrete subgroup of $PSL(2,\mathbb{C})$, which acts on $\mathbb{C}^*$ in the usual way by Möbius transformations, and $\Omega(\Gamma)$ is the set of points on the Riemann sphere where this group 'acts nicely'. More precisely, $\Omega(\Gamma)$ is the set of points $z \in \mathbb{C}^*$ which have some neighborhood U containing no other images of z under the group: $\gamma \cdot z \in U$ for $\gamma \in \Gamma$ implies that γ is the identity. Equivalently, if we define the limit set $\Lambda(\Gamma)$ to be the set of accumulation points of the action of Γ on $\mathbb{C}^*$ (a set about which we will have much more to say later), then $\Omega(\Gamma)$ is just the Riemann sphere with those points removed: $\Omega(\Gamma) = \mathbb{C}^* - \Lambda(\Gamma)$. More specifically, a Schottky group Γ of genus g is a subgroup of $PSL(2,\mathbb{C})$ that is freely generated by g loxodromic[9] elements of $PSL(2,\mathbb{C})$, having as a fundamental domain the exterior of $2g$ closed curves (usually circles), such that each of the g generators of Γ maps one of these boundaries to another in pairs. Intuitively, to obtain a Schottky representation, we can cut the surface along g disjoint closed loops such that it stays in one piece and becomes a sphere with $2g$ holes, flatten it onto the complex plane, and build the Schottky group from the Möbius maps that glue the surface back together along its g seams. A given Riemann surface can be written as a Schottky group in many different ways, depending on the choice of g cycles to cut along. The presentation as a Schottky group is equivalent to the plumbing construction used in (Cho, Collier, and Yin, 2017b). A more detailed review of Schottky uniformization can be found in (Krasnov, 2000; Faulkner, 2013; Barrella et al., 2013). A slightly different approach to calculations in the Schottky coordinates was used in (Gaberdiel, Keller, and Volpato, 2010).

A rather different presentation of the Riemann surface $\mathcal{B}$ is as an algebraic curve. In this case we represent $\mathcal{B}$ as the set of solutions to an equation such as

$$y^n = \prod_{k=1}^{N} \frac{z - u_k}{z - v_k} \,. \tag{4.1}$$

Here, $\mathcal{B}$ is a genus $g = (N-1)(n-1)$ surface, represented as n-fold branched cover over the Riemann sphere parameterized by the z-plane, with $2N$ branch points

a function of τ alone. This dependence, however, involves only the central charge and not any of the other dynamical data of the CFT (such as operator dimensions or structure constants), so will not be important for us here. We will therefore suppress this dependence and simply indicate the dependence on the conformal structure moduli τ.

[9] A loxodromic element γ of $PSL(2,\mathbb{C})$ is one which is conjugate to $\begin{pmatrix} q & 0 \\ 0 & q^{-1} \end{pmatrix}$ for some q with $0 < |q| < 1$.

(u_k, v_k). As the resulting Riemann surface automatically possesses a $\mathbb{Z}_n$ symmetry (usually referred to as replica symmetry) where one permutes the n sheets, one cannot describe a general point in moduli space $\mathcal{M}_g$ using this parameterization. Instead, this equation describes only a $2N-3$ dimensional slice of moduli space, a family of surfaces with an enhanced ($\mathbb{Z}_n$) automorphism group. Except in special cases one cannot map out the full moduli space this way.[10] The advantage of this approach is that the moduli of this surface are easy to describe – they are the locations (u_k, v_k) of the branch points. For physics, this presentation is natural for describing Rényi entropies (or certain correlation functions in orbifold theories), with the branch points u_k and v_k corresponding to the insertion points of twist and anti-twist operators.

Except at genus one, or in special cases with very high symmetry, it is not possible to find an explicit map between the moduli of the algebraic curve and Schottky groups, or to find out whether two Schottky groups represent the same surface, sliced in a different way. However, the problem of finding a Schottky group associated with a particular algebraic curve, sometimes called 'Schottky uniformization', is equivalent to solving a monodromy problem, which we now briefly describe.

To do this, we begin by denoting the locations of branch points (u_k, v_k) as z_i ($i = 1, \ldots, 2N$). We would like to find the map $w(z)$ from the algebraic curve coordinate z to the coordinate w of the complex plane on which the Schottky group acts. But $w(z)$ is not single-valued, because there are many possible values of w related by elements of the Schottky group Γ. However, the Schwarzian derivative $T_c(z) = S(w)(z) = \left(\frac{w''}{w'}\right)' - \frac{1}{2}\left(\frac{w''}{w'}\right)^2$ is single-valued, since the Schottky group consists of Möbius maps. If we take $T_c(z)$ as given, a simple calculation shows that solving $T_c(z) = S(w)(z)$ for w is equivalent to solving the ordinary differential equation

$$\psi''(z) + T_c(z)\psi(z) = 0, \tag{4.2}$$

with $w(z) = \frac{\psi_1(z)}{\psi_2(z)}$ being the ratio of two linearly independent solutions $\psi_{1,2}(z)$ to the ODE.

This is not much use if we know nothing about $T_c(z)$. However, for the Riemann surface (4.1) $T_c(z)$ can be fixed up to a finite number of parameters, by using the fact

[10] One important special case is the genus two moduli space. Every genus two curve is hyperelliptic, so can be represented as a 2-fold cover of the sphere branched over 6 points.

that it is a meromorphic function of z which transforms like a stress-tensor:

$$T_c(z) = \sum_{i=1}^{2N} \left(\frac{1 - n^{-2}}{4(z - z_i)^2} + \frac{\gamma_i}{z - z_i} \right) \tag{4.3}$$

Here, we have assumed that the Schottky group respects the replica symmetry, so $T_c(z)$ is single-valued in z. The double poles are fixed by demanding smoothness in the y coordinate of 4.1, and the γ_i are free parameters, called 'accessory parameters'. It is also constrained by smoothness at infinity, which demands that $T_c(z)$ decays like $\frac{1}{z^4}$. This imposes three constraints on the γ_i, leaving $2N - 3$ free parameters.

It remains to fix these free parameters. To do this, note that if we go around a closed curve on the surface, a solution to the ODE will not usually come back to itself, but undergo monodromy, so the value of w will change by a Möbius map:

$$\begin{pmatrix} \psi_1 \\ \psi_2 \end{pmatrix} \longrightarrow \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} \psi_1 \\ \psi_2 \end{pmatrix}, \quad \text{so that} \quad w \longrightarrow \frac{aw + b}{cw + d} \tag{4.4}$$

The monodromies of the ODE form a representation of the fundamental group of the surface $\pi_1(\mathcal{B})$ by Möbius maps. But in the Schottky representation, not all the closed loops on the surface should take us to a different w, and a different copy of the fundamental domain for Γ: the g special loops that bound the fundamental domain should come back to the same value of w, and so correspond to trivial monodromy of the ODE 4.2 (in fact, the monodromy matrix around these cycles is always minus the identity).

Imposing these trivial monodromy conditions is precisely enough to fix the $2N - 3$ free parameters. Once these parameters are fixed, we may solve 4.2 to find the monodromy around g complementary cycles, which give the generators of the Schottky group Γ. In 4.2 we go into more detail for a specific example, which we will subsequently use for analytic and numerical calculations.

This monodromy problem also appears in computations of the semiclassical limit of Virasoro conformal blocks (Al B Zamolodchikov, 1987), reviewed in (Harlow, Maltz, and Witten, 2011; Hartman, 2013), and described in generality for higher genus blocks in (J. Cardy, Maloney, and Maxfield, 2017).

Schottky representations, handlebodies, and gravity

The Schottky representation has a very natural interpretation from the bulk point of view. To see this, note that the Möbius maps acting on the Riemann sphere can

be extended into a bulk hyperbolic space $\mathbb{H}^3$, where they act as the orientation-preserving isometries. We can therefore extend the action of a Schottky group Γ into this bulk, obtaining a quotient of hyperbolic space with $\mathcal{B}$ as its boundary, $\mathcal{M} = \mathbb{H}^3/\Gamma$. Representing hyperbolic space in the upper half-space model, this can be understood as taking the circles that bound the fundamental domain of the Schottky group and extending them as hemispheres into the bulk, giving a fundamental domain with the hemispheres identified by the generators of Γ.

A CFT partition function on the surface $\mathcal{B}$ can be computed holographically as the bulk gravity path integral over Euclidean geometries whose conformal boundary is $\mathcal{B}$. Semiclassically, we just need to compute the action of a solution to the bulk equations of motion with boundary $\mathcal{B}$, which will depend on the moduli of $\mathcal{B}$. There are an infinite number of bulk solutions, which in general should include the contribution of matter fields, but a particularly simple class of solutions are those without matter fields turned on. Since pure gravity in three dimensions is locally trivial, Einstein's equations then imply that the bulk is locally $\mathbb{H}^3$, which means that it must be a quotient of hyperbolic space. The Schottky group quotients therefore provide a large class of solutions to the bulk problem, which are conjectured to dominate the path integral in pure gravity.

Topologically, the Schottky group quotients are handlebodies, obtained by 'filling in' the surface $\mathcal{B}$ along a choice of g cycles. These contractible cycles are precisely those we chose to cut the surface along to construct the Schottky group, or around which we imposed trivial monodromy. The Schottky group describes the remaining non-contractible cycles, in the sense that it is topologically interpreted as the fundamental group of the handlebody.

Some geometric properties of the bulk can be read off easily from the Schottky group, in particular the lengths of closed geodesics. A closed loop is represented topologically as a conjugacy class in the fundamental group, or equivalently in Γ, and since the eigenvalues of $\gamma \in \Gamma$ are independent of the representative of the conjugacy class, the smaller eigenvalue q_γ of $\gamma \in SL(2,\mathbb{C})$ ($0 < |q_\gamma| < 1$) is naturally associated with a closed curve. Writing $q_\gamma = e^{-\frac{1}{2}(\ell + i\vartheta)}$, ℓ is in fact the length of the closed geodesic, and θ is the amount the geodesic is twisted by (the angle a normal vector rotates by after parallel transport round the curve). Explicitly in terms of the trace, this length is

$$\ell_\gamma = \cosh^{-1}\left[\left|\frac{\mathrm{Tr}\,\gamma}{2}\right|^2 + \left|\left(\frac{\mathrm{Tr}\,\gamma}{2}\right)^2 - 1\right|\right]. \tag{4.5}$$

For more detailed review and discussion, focussing on the Lorentzian versions of these geometries, see (Maxfield, 2015).

Because there are many possible Schottky groups corresponding to the same surface $\mathcal{B}$, we must decide which geometry gives the correct semiclassical bulk dual for given moduli of the boundary surface (even before considering bulk matter fields). The naïve answer to this, and the one that reproduces CFT expectations, is that the handlebody with least action dominates the path integral, so is the dual bulk. It was shown how to compute this action in (Krasnov, 2000), from a particular higher-genus 'Liouville action' (Zograf and Takhtadzhyan, 1988), depending crucially on the IR cutoff imposed on the bulk, and hence on the choice of metric on the boundary surface within the given conformal class. In a metric appropriate for Rényi entropies, flat away from conical singularities at branch points z_i, and with a bulk preserving replica symmetry, the derivative of the action with respect to z_i is proportional to the accessory parameter γ_i (Faulkner, 2013). In a constant curvature metric, for general surfaces, a numerical algorithm to compute the action was given in (Maxfield, Ross, and Way, 2016). As a heuristic, to choose the dominant saddle point, the g shortest cycles of the surface should be filled in.

Given these tools, one can then attempt to construct the the higher genus partition function $Z_g(\tau)$ via a bulk path integral, as a sum over geometries. The handlebodies described above give semi-classical saddle point contributions to this bulk path integral, and the full partition function should be given be a sum over these semi-classical saddles along with a set of loop corrections. The loop corrections to these semi-classical contributions can be computed exactly at genus $g = 1$ and perturbatively at higher genus (see e.g. (Maloney and Witten, 2010; Headrick et al., 2015)). In pure gravity – i.e. in theories with no degrees of freedom aside from the metric – there is some hope that one could compute the higher genus partition function exactly (Witten, 2007). We will be interested in more general theories, which contain scalar fields in addition to the metric. In this case the theory has local bulk degrees of freedom, and there is little hope of an exact computation. Nevertheless, the computations described above give contributions to the partition function of a holographic CFT which will be valid in the semi-classical (large c) limit.

A $\mathbb{Z}_n$ symmetric family of genus $n-1$ surfaces

We now illustrate this general discussion with an example, specifically a family of genus $n-1$ surfaces with an enhanced $\mathbb{Z}_n$ symmetry. As an algebraic curve, this family of Riemann surfaces is given by

$$y^n = \frac{z(z-1)}{z-x} \,. \tag{4.6}$$

This is the $N = 2$ case of (4.1), where we have used $PSL(2,\mathbb{C})$ transformations to put $u_1 = 0$, $u_2 = 1$ and $v_2 = \infty$. The remaining parameter is the cross-ratio x, which is the modulus of this family of Riemann surfaces. In general x can be any complex number, but for simplicity (and for the purposes of applications to Rényi entropies, described below) we will take it to be a real number between zero and one.

To find a Schottky group, or equivalently a bulk geometry, we can now solve the monodromy problem described above. Choosing to preserve the replica symmetry, the most general ansatz for the ODE 4.2 is

$$T_c(z) = \frac{1-n^{-2}}{4}\left(\frac{1}{(z-x)^2}+\frac{1}{z^2}+\frac{1}{(z-1)^2}-\frac{2}{z(z-1)}\right)-\gamma\frac{x(1-x)}{z(z-1)(z-x)}, \tag{4.7}$$

where we have imposed the constraints (which are slightly different, because there is a branch point at infinity), leaving the single accessory parameter γ. To fix this parameter, we must first choose the cycles around which we impose trivial monodromy. For our purposes, it suffices to take the cycles surrounding 0 and x; this gives a loop on the z plane enclosing one zero and one pole of y^n, so remains on the same sheet of the branched cover, forming a closed loop on the surface. There are n of these, one on each sheet, but only $n-1$ of them are in fact independent: the product (in the fundamental group) of the n loops, described by a loop enclosing $z = 0$ and $z = 1$, then moving to the next sheet, and repeating n times, is topologically trivial. Because we have imposed $\mathbb{Z}_n$ symmetry already on our ansatz for $T_c(z)$, imposing trivial monodromy on any one of the sheets is sufficient.

Having chosen the accessory parameters to trivialize the monodromy around these cycles, we would like to read off the Schottky group. To do this, it is convenient to take full advantage of the symmetry of the situation, using the automorphisms of the surface (which are preserved by the handlebody). In the language of the quotient, an (orientation-preserving) isometry of the bulk is represented by an additional element $\gamma \in PSL(2,\mathbb{C})$ (so it is an isometry on the covering space $\mathbb{H}^3$) that commutes with

the group Γ ($\gamma\Gamma = \Gamma\gamma$, so it has a well-defined action on the quotient). Including some such elements, we can form an extended group $\hat{\Gamma}$, of which Γ forms a normal subgroup. The largest possible $\hat{\Gamma}$, including all elements of $\gamma \in PSL(2,\mathbb{C})$ such that $\gamma\Gamma = \Gamma\gamma$ is the *normalizer* $\mathcal{N}(\Gamma)$ of Γ, and the isometry group of the bulk is then $\mathrm{Isom}(\mathcal{M}) \simeq \mathcal{N}(\Gamma)/\Gamma$.

The most obvious extension providing an automorphism is the $\mathbb{Z}_n$ replica symmetry R, represented as the monodromy around a loop containing 0. This comes back to a different sheet, so is not an element of Γ, but we can include it in $\hat{\Gamma}$ as an elliptic Möbius map of order n. From the point of view of the monodromy problem, these elements correspond to monodromy along curves that may not return to the same point on the surface, but go between some point and its image under the isometry. Near $z = 0$, the independent solutions to the ODE 4.2 look like $\psi_\pm(z) \sim z^{\frac{n\pm1}{2n}}$, with corrections forming a power series in z and not affecting the monodromy around zero; choosing these as our basis $\psi_\pm$ of solutions ($w(z) = \frac{\psi_+(z)}{\psi_-(z)}$), the loop around zero, enacting the replica symmetry, acts on the w coordinate as $R : w \mapsto e^{\frac{2\pi i}{n}} w$.

In fact, this family of surfaces automatically has more symmetry, containing an additional $\mathbb{Z}_2$ extending the $\mathbb{Z}_n$ to a dihedral group D_{2n}[11]. From the z coordinate, this can be understood as a map swapping 0 with 1 and x with ∞, $z \mapsto \frac{x}{z}$, along with reversing the order of the sheets of the cover. It is straightforward to check that this leaves the ODE invariant, after transforming ψ as a weight $-\frac{1}{2}$ field. In terms of the monodromy, this extra symmetry is enacted by taking the solutions $\psi_\pm(z)$, following the solution from 0 to ∞, and reading off the coefficients of $z^{\frac{n\pm1}{2n}}$ in these solutions near ∞, giving some order two Möbius map S.

In practice, except for the special case $n = 2$, finding S requires doing the calculation numerically, but we can deduce a lot about it, reducing the unknown parameters from the three numbers specifying a general Möbius map, to just one. Firstly, note that doing S twice corresponds to going round a loop with trivial topology, which implies that S is order two, $S^2 = 1$, which means it is specified by its fixed points. Secondly, without altering the form of R, we can change coordinates by rescaling and rotating in the w plane, and use this to remove one other parameter of freedom. We will use this freedom to set the product of the fixed points of S to be unity, which fixes S to act as $S : w \mapsto \frac{w-\zeta}{\zeta w-1}$ for some (in general complex) ζ.

[11] Even further than this, these surfaces all have another additional $\mathbb{Z}_2$ commuting with this dihedral group, acting as $z \mapsto x\frac{z-1}{z-x}$, which is a hyperelliptic involution of $\mathcal{B}$. We will not make use of this extra symmetry, but in the parameterization used below, it can be included as $w \mapsto \frac{1}{w}$.

Now, the extended group $\hat{\Gamma}$ is generated by just R and S (in fact, it is the free group generated by those elements with the only relations being those given by the orders of the elements: $\hat{\Gamma} = \langle R, S | R^n = S^2 = 1 \rangle$), so we are interested in the one-parameter family of groups generated by one Möbius map of order 2, and one of order n. The actual Schottky group Γ appears as a normal subgroup of this, generated by the loops that actually return to the same point on the surface (but are still non-contractible in the bulk), requiring an even number of S generators to appear, and also for the sheets to map back to themselves, rather than being permuted. The first of these is the element $\gamma_1 = SRSR$, taking a loop round zero by R, then going from 0 to ∞ by S, then a loop round to infinity by R again, and finally back to the starting point at zero, creating a closed loop surrounding zero and infinity, or equivalently x and 1. The remaining generators are similar, but starting on a different sheet, achieved by conjugating with R: $\gamma_k = R^{1-k} SRSR^k$. There are n of these, but they are not all independent, since $\gamma_1 \gamma_2 \cdots \gamma_n = 1$. Any $n - 1$ of these (a number equal to the genus) generate Γ. To relate this to the general discussion of symmetries above, the group of isometries $\hat{\Gamma}/\Gamma$ described by this extension is the dihedral group of order $2n$, since modding out by Γ is equivalent to imposing the additional relation $\gamma_1 = SRSR = 1$, giving the presentation $\hat{\Gamma}/\Gamma = \langle R, S | R^n = S^2 = SRSR = 1 \rangle \equiv D_{2n}$.

This prescribes the family of Schottky groups we are interested in, parameterized by $\zeta = \cos\theta$, though it is important to note that this only describes a Schottky group when ζ is sufficiently close to one (or θ close to zero). An alternative, more geometric parameterization is by q, defined as the smaller eigenvalue of $\gamma_1 = SRSR \in SL(2, \mathbb{C})$ (noting that this is independent of the sign chosen for the matrix representatives of S and R), defined as above so that $q = e^{-\frac{1}{2}(\ell + i\vartheta)}$ gives the length and twist of a curve in the bulk geometry. In the case $n = 2$, $q = \tan^2 \frac{\theta}{2}$ is the usual elliptic nome of the boundary torus, lying in the punctured open unit disc, though for larger n it must be contained in a strictly smaller region.

In practice, to map from any given x to find the corresponding value of ζ (or θ or q), it is sufficient to compute the trace of any of the γ_k (all are equal), from the trace of the monodromy of any loop containing x and 1. It is also possible to solve the monodromy problem perturbatively in small cross-ratio, as described in (Barrella et al., 2013), and in our parametrisation, the result to leading order is $\theta = \frac{\sqrt{x}}{n}(1 + O(x))$. Real x between 0 and 1 corresponds to real θ, or $0 < \zeta < 1$, or $0 < q < 1$. Real negative x (or equivalently $x > 1$) also results in a Fuchsian group Γ, and corresponds to $\zeta > 1$ (but bounded by $\sec \frac{\pi}{n}$ so the group is Schottky). In

this paper we'll use the case $n = 3$ in a number of our examples so it's worth writing down an explicit form for the generators in terms of q. One such parameterization for the $\mathbb{Z}_3$ symmetric Schottky groups yields

$$\langle R, S | R^3 = S^2 = 1 \rangle \cong \langle A, B \rangle \tag{4.8}$$

where

$$A = \begin{pmatrix} -\frac{q\left(q+\sqrt[3]{-1}\right)+1}{\left(-1+(-1)^{2/3}\right)q} & \frac{(-1)^{2/3}(q+1)\left(q+(-1)^{2/3}\right)\sqrt{\frac{(-1)^{2/3}q+1}{q+(-1)^{2/3}}}}{\left(-1+(-1)^{2/3}\right)q} \\ \frac{(q+1)\left(\sqrt[3]{-1}q-1\right)\sqrt{\frac{(-1)^{2/3}q+1}{q+(-1)^{2/3}}}}{\left(-1+(-1)^{2/3}\right)q} & \frac{q\left((-1)^{2/3}q+(-1)^{2/3}+1\right)+(-1)^{2/3}}{\left(-1+(-1)^{2/3}\right)q} \end{pmatrix} \tag{4.9}$$

$$B = \begin{pmatrix} -\frac{q\left(q+\sqrt[3]{-1}\right)+1}{\left(-1+(-1)^{2/3}\right)q} & -\frac{(q+1)\left(\sqrt[3]{-1}q-1\right)\sqrt{\frac{(-1)^{2/3}q+1}{q+(-1)^{2/3}}}}{\left(-1+(-1)^{2/3}\right)q} \\ -\frac{(-1)^{2/3}(q+1)\left(q+(-1)^{2/3}\right)\sqrt{\frac{(-1)^{2/3}q+1}{q+(-1)^{2/3}}}}{\left(-1+(-1)^{2/3}\right)q} & \frac{q\left((-1)^{2/3}q+(-1)^{2/3}+1\right)+(-1)^{2/3}}{\left(-1+(-1)^{2/3}\right)q} \end{pmatrix} \tag{4.10}$$

We can also find different handlebodies for the same surface by trivializing monodromy around some different cycle, but the only possibilities preserving the replica symmetry are much the same, the most obvious being to take the loop to surround x and 1 rather than 0 and x. There is a phase transition between the two corresponding handlebodies at $\mathrm{Re}(x) = \frac{1}{2}$ (Headrick, 2010).

Relationship with Rényi entropies

The bulk computation of higher genus partition functions can be applied to the computation of Rényi entropies in holographic CFTs (see (Calabrese and J. Cardy, 2009) for a review). For a density matrix ρ, the nth Rényi entropy is defined as

$$S_n = \frac{1}{1-n} \log \mathrm{Tr}\, \rho^n. \tag{4.11}$$

In the $n \to 1$ limit this becomes the von Neumann entropy $S = -\mathrm{Tr}(\rho \log \rho)$. In order to probe the spatial entanglement structure of the theory, we can take ρ to be the reduced density matrix for a spatial region A in the vacuum state. Then ρ is defined by a Euclidean path integral on the sphere (with cuts introduced at A), and the Rényi entropy may be computed by gluing n copies of this sphere together along these cuts. Explicitly, we have

$$S_n = \frac{1}{1-n} \log \frac{Z_n}{Z_1^n}. \tag{4.12}$$

Here Z_n is the partition function on a manifold M_n, which is the n-fold branched cover defined by gluing n copies of the original spacetime manifold along A, and the normalization constant Z_1 is the sphere partition function. If A consists of N disjoint intervals, this is precisely the $n-$fold cover of the sphere branched over $2N$ points (the endpoints) described above.

We will focus on the case where A consists of two disjoint intervals $[u_1, v_1]$ and $[u_2, v_2]$. Then the conformal structure of M_n is completely determined by the cross ratio

$$x = \frac{(u_1 - v_1)(u_2 - v_2)}{(u_1 - u_2)(v_1 - v_2)}. \tag{4.13}$$

As this cross-ratio is varied, we sweep out a one (real) dimensional slice of the moduli space $\mathcal{M}_{n-1}$ of genus $n-1$ dimensional Riemann surfaces. This is precisely the case described in the previous subsection. There are two handlebodies which compute the pure-gravity contribution to the Rényi entropies, which exchange dominance at the point $x = 1/2$. When $n = 2$ these saddles are precisely the thermal AdS and Euclidean BTZ black hole solutions, and the phase transition at $x = 1/2$ is the usual Hawking-Page phase transition. The scalar instabilities we describe in this paper will occur for $n > 2$ when the theory has an operator of dimension Δ which is sufficiently light. In particular, we will find that there are two new phase transitions as x is varied, one at $x = x_c(\Delta) < 1/2$ and one at $x = 1 - x_c(\Delta)$, where these two handlebodies will become unstable to the formation of scalar hair.

4.3 The phase transition from CFT

In this section we will make a CFT argument for the instability, by considering the contribution of a generalized free field – the boundary avatar of a free bulk scalar field – to the higher genus partition function of a CFT with large central charge. The result is that if the corresponding operator is sufficiently light, then the generalized free partition function will diverge somewhere in the interior of moduli space. This signals that the free approximation has broken down, so interactions become important, and the partition function will undergo a phase transition. We give an analytic lower bound on the critical dimension Δ_c in terms of the Schottky moduli of the surface. As the field becomes lighter, the phase transition will occur closer to the boundary of moduli space; in particular, for a sufficiently light field the corresponding bulk phase transition will occur before the usual "Hawking-Page" transition where (locally Einstein) bulk saddles are interchanged. In the bulk, this would be interpreted as the condensation of a bulk scalar field in a

handlebody background. The discussion in this section is a refinement of the arguments presented in (Belin, Keller, and Zadeh, 2017).

The partition function and conformal blocks

A higher genus partition function can, at least in principle, be computed from the basic dynamical data of the CFT, namely the spectrum of dimensions and spins (Δ_i, s_i) of primary operators, along with their three-point coefficients C_{ijk}. To do this, we can insert a complete set of states on a handle of the surface to reduce the computation to sum over two-point functions on a surface one genus lower, and repeat this (along with use of the OPE) until the computation has been reduced to three-point functions on spheres. A complete decomposition like this can be understood by cutting up the surface into pairs of pants: any genus $g \geq 2$ surface can be decomposed (in many ways) into $2(g-1)$ pairs of pants, joined along a total of $3(g-1)$ cuffs. Along each of these cuffs, we can insert a complete set of states and, by the state-operator correspondence, the amplitude between three states defined by the path integral on the pair of pants is determined by a three-point coefficient.

Along with inserting complete sets of states in this way, we can use the fact that the states are arranged in multiplets of the Virasoro algebra, by summing up all contributions from a given multiplet appearing in the sums over states. The resulting object, collecting the contributions from a given primary on each of the $3(g-1)$ cuffs along with all their descendants, is a higher genus conformal block $\mathcal{F}$. This is determined by kinematics alone, depending only on the scaling dimensions and spins of the primaries chosen, the central charge, and the moduli of the surface[12]. In the partition function this conformal block will be multiplied by the product of $2(g-1)$ OPE coefficients corresponding to the primaries.

Summing over all possible choices of primaries, the result is a general expression for the partition function of the form

$$Z_g(\tau) = \sum_{\{i\}} C_{\{i\}}^{2(g-1)} \mathcal{F}(\{\Delta_i\}, c; \tau) \tag{4.14}$$

Here the sum is over all choices of the $3(g-1)$ primary operators, and $C_{\{i\}}^{2(g-1)}$ denotes the product of $2(g-1)$ OPE coefficients corresponding to the primary operators propagating down the legs of each pair of pants. This expression may

[12] The blocks also factorize into the product of holomorphic and antiholomorphic blocks, though we will not explicitly use this fact here.

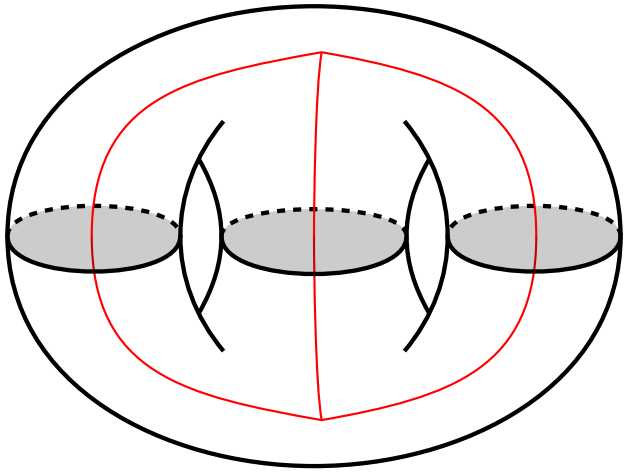

Figure 4.1: A genus 2 surface, which is cut into two pairs of pants glued together along the three black circles. Along each of the three circles we can insert a projection onto the descendants of a primary of dimension Δ_i ($i = 1, 2, 3$) to obtain the block $\mathcal{F}(\{\Delta_i\}, c; \tau)$. The dual handlebody is found by 'filling in' the surface, as indicated by the shaded disks. The block $\mathcal{F}(\{\Delta_i\}, c; \tau)$ can be computed in the bulk in a semi-classical approximation, valid in the limit $1 \ll \Delta_i \ll c$, by computing the action of the network of bulk geodesics indicated in red.

look very different for different pair-of-pants decompositions of the surface, but the result must be equal whichever decompositions is chosen; this is the statement of higher-genus crossing symmetry, which can be exploited to constrain CFT data (J. Cardy, Maloney, and Maxfield, 2017; Cho, Collier, and Yin, 2017a; Keller, Mathys, and Zadeh, 2017).

In the case $g = 2$, there are two possible distinct types of decomposition into pairs of pants, depending on whether we choose to insert a complete set of states on a cycle dividing the surface into two pieces. Assuming we do not, the decomposition looks like two pairs of pants joined to one another along each of their three cuffs, as illustrated in 4.1, and $C^{2(g-1)}_{\{i\}}$ is just C^2_{ijk}, where i, j, k denote the primaries chosen on each of these seams.

In general, it is rather difficult to compute 4.14 explicitly, and even the conformal blocks cannot be calculated exactly in closed form. It is possible to calculate perturbatively in moduli of the surface, or using recursion relations exploiting the structure of degenerate representations (Cho, Collier, and Yin, 2017a; Cho, Collier, and Yin, 2017b), or in various semiclassical limits (J. Cardy, Maloney, and Maxfield, 2017; Kraus et al., 2017). In particular, one needs to choose a conformal frame, and to account carefully for the way in which the surface is glued together from its

constituents. Fortunately we will not need to work with this expression in generality, only requiring the blocks in a particular 'semiclassical global' limit.

Semiclassical global limit of higher genus blocks

We will require the blocks in a limit of large central charge, where the dimensions of exchanged operators are large also large (with ratios between different Δ_i fixed in the limit), but small compared to c: $1 \ll \Delta_i \ll c$. This limit has a dual holographic description in terms of semiclassical gravity coupled to particles in a probe limit, for which only the global conformal $\mathfrak{sl}(2)$ subalgebra of the Virasoro algebra is important; hence the name 'semiclassical global' blocks.

In this limit the blocks simplify, becoming

$$\mathcal{F}(\{\Delta_i\}, c; \tau) = e^{-cS_0(\tau)-\Delta S_1(\{\Delta_i/\Delta_j\};\tau)+O(1,\Delta^2/c)}, \tag{4.15}$$

The functions S_0 and S_1, depend on the moduli and (in the case of S_1) on the ratios of conformal dimensions, and have semi-classical gravity interpretations which we will describe below. The fact that the blocks exponentiate in this limit large c limit is most well known in the case of four-point functions (A. B. Zamolodchikov, 1984), but has not been rigorously proven. It is, however, physically well-motivated, for example by considering a semiclassical limit of Liouville theory (Harlow, Maltz, and Witten, 2011).

First, S_0 can be interpreted as the semiclassical vacuum block, i.e. the block for which all operators are taken to be the identity. It is equal to the on-shell action of the handlebody where each of the cycles in the pair-of-pants decomposition are chosen to be contractible.[13] This depends on a choice of conformal frame (the conformal anomaly precisely takes the form of a shift in S_0), and in general can only be computed exactly by numerics (Hartman, 2013; Faulkner, 2013; Maxfield, Ross, and Way, 2016). The frame-independent information contained in S_0 is the difference between its value in different channels, and this data is required to impose crossing symmetry. Luckily, for our purposes we will only need to express the partition function in one channel, so we can entirely disregard this piece.

For us, the more important contribution is S_1, encoding the dependence on the dimensions. The important point is that in the limit $c \to \infty$ with fixed Δ_i the con-

[13] This only depends on the choice of g cycles, not the full decomposition, which is consistent because of the fusion rules of the identity: choosing the identity module along g cycles is enough to imply that the identity must be present in the other $2g-3$.

tribution of the Virasoro descendants is unimportant after we factor out the vacuum block contribution S_0. So the block reduces to a 'global' block, where only the L_{-1} descendants are kept (Cho, Collier, and Yin, 2017b)[14]. The gravitational interpretation is clear: as $c \to \infty$, the backreaction from the matter and the loop corrections from the graviton can be ignored, and we need only the classical background action.

This global block is still tricky to compute at higher genus, but if we further assume that the internal dimensions Δ_i are large (but still much smaller than c), it simplifies to a 'semiclassical global' block. This can be determined by considering a network of geodesics in the handlebody spacetime, determined by the pair-of-pants decomposition. Specifically, for each pair of pants, assign a trivalent vertex in the bulk, and join these vertices by a geodesic for each seam joining the pairs of pants. This geodesic is interpreted as the worldline of a particle of mass $m_i \sim \Delta_i$, determined by the dimension of the primary operator assigned to the corresponding cuff of the pants decomposition, and is assigned an action $\Delta_i \ell_i$, where ℓ_i is the length of that geodesic segment. We finally must specify the bulk locations of the vertices; these are chosen to minimize the total particle action $\sum_i \Delta_i \ell_i$. The value of $e^{-\min \sum_i \Delta_i \ell_i}$ at this minimum obeys the semiclassical limit of the Casimir equations for the global conformal blocks, as shown in (Kraus et al., 2017), so reproduces the correct moduli dependence of the blocks.[15] This limit can be used as a starting point for a systematic perturbative expansion for the blocks, developed in terms of worldline quantum field theory coupled to gravity in (Maxfield, 2017).

The only thing that remains to fix is the overall normalization. The geodesic network prescription comes with an unambiguous normalization, but rather than being the canonical one, where we multiply by the appropriate OPE coefficients to find the contribution to the partition function, it comes with some nontrivial OPE coefficients $\hat{C}(\Delta_i)$ built in, depending on the dimensions of the operators meeting at each vertex. To compute these, consider taking the pinching limit in which all cuffs of the pairs of pants become small, suppressing the descendants so only the product of primary three-point functions remains. The function $\hat{C}(\Delta_i)$ can therefore be computed by using a geodesic approximation to a three-point function, with three geodesics going

[14] What constitutes a global descendant is a little ambiguous for higher-genus surfaces, since it is not invariant under general conformal transformations. The statement here requires a Schottky, or plumbing frame, for which all transition maps are Möbius maps.

[15] For some values of the dimensions and moduli, the action can be minimized when one of more of the geodesics shrink down to zero size, in which case the block is given instead by some complexified saddle point. From the bulk point of view, this happens when the leading order amplitude in the large Δ limit comes from double-trace contributions, rather than the case we would like to consider, where these are exponentially suppressed in Δ relative to the single traces (Maxfield, 2017).

from the boundary of AdS and meeting at a trivalent vertex (Chang and Lin, 2016):

$$\hat{C}(\Delta_i) = e^{\mathcal{P}(\Delta_i)}, \text{ where} \tag{4.16}$$
$$\mathcal{P}(\Delta_i) = \frac{1}{2}\Delta_1 \log\left[\frac{(\Delta_1+\Delta_2-\Delta_3)(\Delta_1+\Delta_3-\Delta_2)}{\Delta_2+\Delta_3-\Delta_1}\right] + (2 \text{ permutations})$$
$$+ \tfrac{1}{2}\left(\textstyle\sum_i \Delta_i\right)\left(\log \textstyle\sum_i \Delta_i - \log 4\right) - \textstyle\sum_i \Delta_i \log \Delta_i$$

Alternatively, the same result can be obtained from an appropriate limit of the DOZZ formula (Dorn and Otto, 1994; Alexander B. Zamolodchikov and Alexei B. Zamolodchikov, 1996). The function $\mathcal{P}$ is homogeneous of degree one in the dimensions, so gives a contribution scaling linearly with dimension in the exponential, as required. We will make particular use of the special case where all dimensions Δ_i are equal to Δ_p, for which $\mathcal{P} = -\frac{3}{2}\log(\frac{4}{3})\Delta_p$.

In the end, this gives the expressions for the blocks that we will use, applying in the limit $1 \ll \Delta_i \ll c$:

$$\mathcal{F}(\Delta_i, c) \sim \exp\left[-cS_0 - \min\left\{\sum_{\text{edges}} \Delta_i \ell_i\right\} - \sum_{\text{vertices}} \mathcal{P}(\Delta_i)\right] \tag{4.17}$$

We will apply this result in the specific case of the $\mathbb{Z}_3$ symmetric, genus 2 handlebodies, with time-reflection symmetry, corresponding to the $n = 2$ version of the example in 4.2, with x (or θ) real. The relevant geodesic network for the channel of interest is shown in red, in 4.1. Furthermore we will take the dimensions of the three internal operators to be equal, $\Delta_1 = \Delta_2 = \Delta_3 = \Delta_p$. In this case, finding the location of the vertices is straightforward, since they are fixed completely by symmetry, absolving us of the need to solve the minimization problem. It is now a simple exercise in hyperbolic geometry to work out the length of the geodesics connecting the two vertices, finding $\ell = \log\left(\cot^2\frac{\theta}{4}\right)$. This, along with $\mathcal{P} = -\frac{3}{2}\log(\frac{4}{3})\Delta_p$ as computed above, gives the result we will need for the block:

$$\mathcal{F}_{g=2}(\Delta_p, c) \sim e^{-cS_0}\left(\frac{4}{3}\tan^2\frac{\theta}{4}\right)^{3\Delta_p} \tag{4.18}$$

The intuition behind the derivation of this expression relies on the operators in the internal channels being single trace operators, corresponding to single particle states, in a theory with semiclassical gravity dual. But because the blocks are kinematic objects, these restrictions are not required to apply the formula. We will use it in the case where the internal operators are highly composite multi-traces built from a primary of small dimension, for which the intuition behind the semiclassical blocks certainly does not hold.

Applying the blocks to generalized free fields

Consider a scalar $\mathcal{O}$ of dimension Δ, dual to a weakly interacting bulk field. As long as these interactions are unimportant, we can treat $\mathcal{O}$ as a generalized free field, which means that we can sensibly talk about composite 'multi-trace' operators built from products of $\mathcal{O}$ and derivatives, $:\partial^{\#}\mathcal{O}\cdots\partial^{\#}\mathcal{O}:$. In the generalized free approximation, the dimensions of these products simply add, and they have vanishing connected correlation functions, so the correlators can be computed by Wick contractions.

Now, let's try to compute the genus two partition function using the conformal block decomposition, accounting for such a free bulk scalar field. It is a slightly tricky prospect accounting for all the possible multi-trace exchanges, so we will make a slightly crude approximation, taking the contribution only of primary operators $:\mathcal{O}^K:$ without derivatives, of dimension $K\Delta$, and also taking the same operator to propagate in all three legs. This gives us a lower bound for the partition function, since the OPE coefficients are real and the blocks are positive[16]:

$$Z_{\text{GF}} \geq \sum_K C_{KKK}^2 \mathcal{F}_{g=2}(K\Delta, c) \tag{4.19}$$

The OPE coefficients appearing can be computed from the combinatorics of the Wick contractions (Belin, Keller, and Zadeh, 2017), and for three identical operators $:\mathcal{O}^K:$, in the limit of large K, the result is

$$C_{KKK} \sim 2^{3K/2}. \tag{4.20}$$

Putting this together with 4.18 giving the blocks in the appropriate limit, the terms in the sum for large K look like

$$C_{KKK}^2 \mathcal{F}_{g=2}(K\Delta, c) \sim 2^{3K} \left(\frac{4}{3}\tan^2\frac{\theta}{4}\right)^{3K\Delta}. \tag{4.21}$$

But now, if Δ is sufficiently small, these terms grow exponentially, causing the partition function to diverge! We can therefore put a bound on the critical dimension Δ_c at which $Z_{\text{GF}}(\Delta)$ diverges:

$$\Delta_c \geq \frac{\log 2}{\log\left(\frac{3}{4}\cot^2\frac{\theta}{4}\right)}. \tag{4.22}$$

[16] To prove this, note that when x is real, the surface can be constructed by gluing a pair of pants directly to a reflected version of itself. The path integral on the pair of pants defines a state on three copies of the CFT, and the block (times OPE coefficients) is the expectation value of a projection (a positive operator) in this state, which is positive by unitarity.

We do not expect this to be exact, since we have dropped the contribution of so many operators, but we will see later that it becomes asymptotically equal to the correct value at small θ, corresponding to small x. For cross-ratio $x = \frac{1}{2}$, numerically solving the monodromy problem described around 4.2, we find the corresponding value $\theta = .55128$. This gives the bound $\Delta_c \geq 0.189219$ on the critical dimension (accurate to the number of quoted decimal places), in agreement with the analysis of (Belin, Keller, and Zadeh, 2017).

A partition function should be well-defined for any surface, so it may seem puzzling to get a divergent answer. The resolution is that the partition function is not truly divergent, but our approximations on the spectrum and OPE coefficients do not apply when K is parametrically large. Even if we do not give a potential to the bulk field, it interacts through gravity, so the approximation of computing OPE coefficients of multi-trace operators by Wick contractions will cease to apply when K is of order $\sqrt{c}$, though it could break down sooner if other interactions become important at a lower energy scale. When we pass the critical dimension, the sum over blocks will shift from being dominated by the vacuum, to being dominated by the multi-particle states at a scale set by the interactions. This signals a second order phase transition, which we will explain from the bulk as condensation of the scalar field.

4.4 The bulk instability

We have argued in Section 3 that the contribution of a free scalar to the genus two partition function will diverge for sufficiently small conformal dimension, $\Delta < \Delta_c$. This divergence comes from the contribution of multi-trace states which are dual in the bulk to states with large particle number. It is therefore natural to expect that the divergence signals an instability where the bulk scalar field condenses to form a new solution with a nonzero classical value. This implies the existence of a second-order phase transition, below which the semiclassical bulk path integral is dominated by a new classical solution of the (nonlinear) bulk equations of motion: a 'hairy handlebody'.

The new classical solution will depend on the details of the theory, and in particular the interactions of the bulk field. These interactions give anomalous dimensions and couplings to the multi-trace operators in the theory, which become important above the scale of the interactions.[17] In particular, they will modify the asymptotic

[17] This scale could be the Planck scale for a free scalar minimally coupled to gravity, the AdS scale for a strongly coupled bulk field, or some intermediate energy scale such as the string scale.

behaviour of the sum described in section 3 in such a way as to cure the divergence. The result is that the partition function will have some non-universal contribution at the interaction scale of the bulk field.

While the full nonlinear solution depends on details of the theory, the onset of the instability does not, and is sensitive only to the background geometry and the mass of the scalar field. It is characterized by the appearance at $\Delta = \Delta_c$ of a zero mode, a nonzero solution of the linearized bulk wave equation with source-free boundary conditions, which corresponds to a flat direction in the path integral. In this section, we will show that such a zero mode exists in quite general circumstances, and characterize the critical dimension Δ_c in terms of the bulk geometry.

The zero mode and spectral theory

In a d-dimensional holographic CFT, a single-trace scalar operator $\mathcal{O}$ of dimension Δ is dual to a bulk scalar field ϕ of mass $m^2 = \Delta(\Delta - d)$. The linearized bulk equation of motion $(\nabla^2 - m^2)\phi = 0$ has two linearly independent solutions with different asymptotic behaviour near the boundary:

$$\phi(x, z) \sim J(x) z^{d-\Delta} + \langle \mathcal{O}(x) \rangle z^{\Delta} \tag{4.23}$$

Here $J(x)$ is a source for the operator $\mathcal{O}$ in the CFT, which is fixed as a boundary condition, and the expectation value $\langle \mathcal{O}(x) \rangle$, in most circumstances determined uniquely by the boundary condition J and regularity, describes the 'response' of the scalar field in the presence of the linearized source J. We will be most interested in relevant operators $\Delta < d$, corresponding to masses which are naïvely tachyonic, but above the Breitenlohner-Freedman (BF) bound, $-\frac{d^2}{4} < m^2 < 0$. In particular, we recall that for $-\frac{d^2}{4} < m^2 < -\frac{d^2}{4} + 1$, there are two possible choices of boundary condition for the scalar field ϕ with unitary duals (Klebanov and Witten, 1999). These two different boundary conditions correspond simply to a choice of which of the two boundary behaviours in 4.23 one chooses to view as a source, and which as a response.

This linearized Klein-Gordon equation suffices to find leading order correlation functions (ignoring backreaction and other interactions), not just on pure AdS, but any asymptotically AdS geometry $\mathcal{M}$ obtained altering the boundary geometry or sourcing other fields, as long as $\phi = 0$ on the background.[18] This corner of

[18] In general, ϕ could have couplings to curvature or other nonzero fields which modify this linearized equation, but we will be interested primarily in geometries which are locally AdS. Thus all of these couplings may be incorporated into an effective bulk mass of the scalar field.

AdS/CFT therefore reduces to the theory of the Laplacian on the manifold $\mathcal{M}$. Even for geometries that are locally AdS, this spectral theory can be rich and interesting, and we will import some ideas and results from the mathematics literature and explore the physical consequences.

Our key result is that, even in the absence of a source for the operator $\mathcal{O}$, it is possible for ϕ to spontaneously acquire a nonzero classical expectation value. The second-order transition to this behaviour occurs when there is a nonzero solution of the bulk wave equation

$$\left(\nabla^2 - \Delta(\Delta - d)\right)\phi = 0 \tag{4.24}$$

with vanishing source $J = 0$. In other words, as we vary the bulk solution $\mathcal{M}$ (or the dimension Δ), the solution will become unstable when there is an eigenfunction of the Laplacian with boundary condition $J = 0$ and eigenvalue $\Delta(\Delta - d)$. For a given geometry $\mathcal{M}$, we call the largest dimension for which this occurs the critical dimension Δ_c. Reducing Δ further, this eigenfunction becomes a mode which decreases the action of the solution, so a given geometry is unstable to condensation of a scalar with $\Delta < \Delta_c$.

It is easy to see that the instability cannot happen if $\Delta > d$. In particular, for a scalar field with $(\nabla^2 - m^2)\phi = 0$ we can use the standard argument for negativity of the Laplacian:

$$0 \leq \int_{\mathcal{M}} (\nabla\phi)^2 = -\int_{\mathcal{M}} \phi\nabla^2\phi = -m^2 \int_{\mathcal{M}} \phi^2 \implies m^2 < 0 \tag{4.25}$$

Here we have integrated by parts, and used the fact that the fast fall-off conditions ($J = 0$ and $\Delta > \frac{d}{2}$) imply that all integrals converge and boundary terms vanish. This instability is therefore ruled out for an irrelevant operator, but not immediately excluded for relevant operators, for which the boundary terms do not automatically vanish. We will see that such instabilities do occur, and are in fact quite generic.

Before giving our first characterization of the onset of the instability in terms of spectral theory, we should first clarify some mathematical terminology. We are seeking a solution ϕ of the equation $\nabla^2\phi = \Delta(\Delta - d)\phi$, with boundary conditions $J = 0$. In the mathematics literature, this would be called an eigenfunction of the Laplacian only if $\Delta > \frac{d}{2}$, since in this case the boundary condition $J = 0$ is equivalent to demanding square-integrability of the eigenfunction: $\phi \in L^2(\mathcal{M})$. We will also be interested in the alternate quantization of the scalar field, corresponding to operators in the range $0 < \Delta < \frac{d}{2}$, where the boundary condition is imposed on

the slowly-decaying mode. The dimension Δ then corresponds to a *resonance* of the Laplacian, which is defined as a pole of the resolvent operator $R_\Delta = (\Delta(\Delta-d)-\nabla^2)^{-1}$, analytically continued in Δ. The resolvent is essentially the bulk Green's function (bulk-to-bulk propagator) G_Δ on $\mathcal{M}$; to compute the action of the resolvent on a function, integrate it against G_Δ, which satisfies the bulk wave equation with delta-function source:

$$R_\Delta[\phi](y) = \int_{\mathcal{M}} d^{d+1}y' G_\Delta(y,y')\phi(y'), \quad (\Delta(\Delta-d)-\nabla_y^2)G_\Delta(y,y') = \delta(y,y') \tag{4.26}$$

The critical dimension Δ_c will therefore show up as a pole in the bulk Green's function G_Δ.

Another way to characterize the critical dimension Δ_c, is to note that the determinant of the resolvent $\det(\Delta(\Delta-d)-\nabla^2)^{-1}$ is precisely the square of the one-loop partition function of ϕ. One can therefore find Δ_c by looking for a divergence in the one-loop contribution of a scalar field ϕ on the background $\mathcal{M}$.

As a final characterisation of Δ_c, we can consider the linear response problem of turning on some small source $J(x)$, solving the bulk wave equation with the corresponding boundary condition, and reading off the response $\langle O\rangle_J$. At generic values of Δ, this problem will have a unique solution, so defines a linear map $S_\Delta : J \mapsto \langle O\rangle_J$ between functions on the boundary $\mathcal{B}$, known in the mathematics literature as the scattering matrix. If we tune Δ to the critical dimension Δ_c, however, there is an ambiguity, as we can always add a multiple of the zero mode to the solution. The zero mode therefore also shows up as a pole of S_Δ, a *scattering pole*.[19] In the same way that the resolvent is related to the bulk Green's function, the scattering matrix is related to the CFT two-point function $\langle O(x_1)O(x_2)\rangle_{\mathcal{B}}$ in the relevant background. This will also diverge as a function of Δ as the critical dimension is approached (the familiar divergence in susceptibility at a second-order phase transition), with a pole at Δ_c, signalling the breakdown of the linearized bulk theory when the scalar becomes unstable.

So far we have been quite general. We will now focus on the case of three dimensions, where we can make more concrete statements about Δ_c.

[19] Scattering poles do not coincide with the resonances for two reasons. The first is that there are also zeros of the scattering matrix, corresponding to solutions with a source but zero response, which may cancel a pole, giving a resonance without corresponding pole in S_Δ. Secondly, the scattering matrix has extra poles at half-integer values of Δ, even in pure hyperbolic space, related to the logarithms that appear in the boundary expansion 4.23 when the asymptotic powers differ by an integer, requiring additional counterterms.

Locally hyperbolic spaces

Let us now consider the case where the bulk geometry $\mathcal{M}$ is a locally hyperbolic space of the form $\mathcal{M} = \mathbb{H}^3/\Gamma$. We will be primarily interested in the case where $\mathcal{M}$ is handlebody, so we will take Γ to be a Schottky group of genus $g > 1$. In fact, the results of this section will apply in greater generality, to non-handlebodies, to some geometries containing conical defects, as well as to hyperbolic manifolds of general dimension.[20]

We will consider a bulk scalar propagating on this geometry, and characterize the relevant spectral theory in terms of properties of the quotient group Γ. We will only motivate and explain the results here, referring to the appropriate mathematics literature for more details, precise statements, and proofs.

Consider first computation of the bulk two-point function of ϕ in the geometry $\mathcal{M} = \mathbb{H}^3/\Gamma$. This can be computed using the method of images, by starting with the two-point function in $\mathbb{H}^3$ and summing over all elements of the group, corresponding to sources at all image points. The result is

$$G_\Delta^{\mathcal{M}}(y, y') = \sum_{\gamma\in\Gamma} G_\Delta^{\mathbb{H}^3}(y, \gamma\cdot y') = -\frac{1}{2\pi}\sum_{\gamma\in\Gamma}\frac{e^{-\Delta d(y,\gamma\cdot y')}}{1-e^{-2d(y,\gamma\cdot y')}} \tag{4.27}$$

where $d(y, y')$ is the geodesic distance between the points y and y', with respect to the $\mathbb{H}^3$ metric. Formally, this gives a function invariant under the group Γ, hence well defined on $\mathcal{M}$, and solves the Klein-Gordon equation with the appropriate source. However, this function will not be well-defined if the sum over images fails to converge. In particular, if the number of image points with $d(y, \gamma\cdot y')$ less than some distance d grows rapidly enough as $d \to \infty$ (for some fixed y, y'), then the sum will diverge. More specifically, if the number of image points with $d(y, \gamma\cdot y') < d$ grows like $e^{\delta d}$, then the sum will diverge for $\Delta < \delta$. In fact, this is always the case for some $\delta > 0$, as stated in the following result of Sullivan (Sullivan, 1979):

Theorem 2 *The series*

$$G_\Delta^{\mathcal{M}}(y, y') = -\frac{1}{2\pi}\sum_{\gamma\in\Gamma}\frac{e^{-\Delta d(y,\gamma\cdot y')}}{1-e^{-2d(y,\gamma\cdot y')}} \tag{4.28}$$

[20] The technical assumptions required are only that Γ is not elementary, which excludes a few simple cases, most notably the cyclic groups corresponding to the Euclidean BTZ geometry, and that it is geometrically finite, which is true in physically relevant cases and in particular for the Schottky groups with genus $g > 1$.

converges in the right half-plane $\mathrm{Re}\,\Delta > \delta$*, where* $\delta > 0$*, the* exponent of convergence *of* Γ*, is the location of the first resonance of* $\mathbb{H}^3/\Gamma$*. The Green's function* $G_\Delta^{\mathcal{M}}(y, y')$ *(analytically continued in* Δ*) has a pole at* $\Delta = \delta$*, and the residue of that pole is given by*

$$\operatorname*{Res}_{\Delta\to\delta} G_\Delta^{\mathcal{M}}(y, y') \propto \phi_0(y)\phi_0(y'), \tag{4.29}$$

where ϕ_0 *is the zero mode function, the solution of the free bulk wave equation with source-free boundary conditions.*

As described in the previous section, this pole in the Green's function, the resonance, signals the onset of an instability. Thus the critical dimension Δ_c equals the exponent of convergence δ. We emphasize that δ is strictly positive given our assumptions on $\mathcal{M}$, which implies that any handlebody of genus greater than one will be unstable if there is a sufficiently light operator in the spectrum.

We may also compute the CFT two-point function of $\mathcal{O}$ in this background, by taking the limit of the bulk Green's function as the points approach the boundary. The exact result for the two point function will depend on the conformal frame, which corresponds to a choice of regulator as we take the points to the boundary, but the convergence properties of the sum over images will be insensitive to this choice. We can write a general metric on the boundary as $ds^2 = e^{2\sigma(w)}dwd\bar{w}$, where w is the complex coordinate on which Γ acts by Möbius maps. The conformal factor σ is defined on the regular set Ω of Γ, and defines a metric on the quotient manifold $\mathcal{B} = \Omega/\Gamma$ under the condition $\sigma(\gamma(w)) = \sigma(w) - \log|\gamma'(w)|$ for all Möbius maps $\gamma \in \Gamma$. The bulk computation of the two-point function on $\mathcal{B}$ gives a sum over images:

$$\langle \mathcal{O}(w)\mathcal{O}(w')\rangle_{\mathcal{B}} = e^{-\Delta\sigma(w)}e^{-\Delta\sigma(w')}\sum_{\gamma\in\Gamma}\frac{|\gamma'(w)|^\Delta}{|\gamma(w) - w'|^{2\Delta}}. \tag{4.30}$$

In the summation on the right hand side, the denominator corresponds to the two-point function on the plane, and the numerator is the conformal factor appropriate for each image. Once again, this sum converges in the right half-plane $\mathrm{Re}\,\Delta > \delta$, and the divergence in the two-point function signals the onset of a second-order phase transition. In the mathematical literature, this CFT two-point function appears as the kernel of the scattering matrix (S. Patterson, 1989).

This sum can be used to gain some intuition about the relationship between the exponent of convergence δ and the geometry of the group Γ. The first thing to note is that the tail of the sum, which controls the divergence, is closely related to

the limit set Λ of the group Γ. The limit set is the set of points where the images $\gamma(w)$ accumulate, for any starting point w. More precisely, a point is in Λ if every neighbourhood of that point contains infinitely many of the images $\gamma(w)$. These are the places where the quotient by Γ acts 'badly', which we must remove to form the regular set $\Omega = \mathbb{C}^* - \Lambda$, so we obtain the nice quotient space $\mathcal{B} = \Omega/\Gamma$. The tail of the sum is controlled by the limit set, since only a finite number of terms in the sum will lie outside of any arbitrarily small neighbourhood of Λ. In the simple case of BTZ, Γ is the cyclic group consisting of the maps $\gamma_n(w) = q^{2n}w$ for $n \in \mathbb{Z}$, and Λ consists of the two points 0 and ∞. In most other cases, however, Λ is much more complicated.

For any limit point (that is, element of the limit set), there is a sequence of images of our starting point w that approach it, say $\gamma_n(w)$ for some $\gamma_n \in \Gamma$ (which are independent of w). As n increases, γ_n will usually be a longer and longer word built out of the generators. For the images to tend to some limit, the γ_n must eventually start with the same string of generators, because if they don't, they would map w to places that are separated by some finite distance: as the sequence γ_n goes on, the words built out of the generators get longer and longer, and only change later and later on in the string. More precisely, the kth letter of the word γ_n is constant after some sufficiently large n, for any k. For each limit point, we can in this way construct a *unique* semi-infinite word built from the generators of the group, a sort of decimal expansion, but using Möbius maps instead of digits. Such words are in one-to-one correspondence with elements of Λ. For $g \geq 2$, this set is not only infinite, but uncountable. The 'rational numbers' in the analogy with decimals consist of strings of generators that eventually repeat, and are in one-to-one correspondence with the primitive elements $\mathcal{P}$ of the group Γ, that is, elements that cannot be written as γ^n for any $n > 1$ (excepting the identity), corresponding to the attractive fixed point of that element.

The resulting set Λ, which controls the tail of the sum over the group, has a rich and beautiful fractal structure. For Fuchsian groups, generated by matrices with real entries, the limit set is a subset of the real line, and closely resembles (indeed, topologically, is homeomorphic to) a Cantor set. Allowing more general Schottky groups, the limit set moves into the complex plane, forming a twisting, intricate, self-similar pattern. Several examples arising when investigating the 3rd Rényi entropy of two intervals are illustrated in 4.2. For many more images of limit sets, and a playful semi-popular account of the mathematics involved, we encourage a

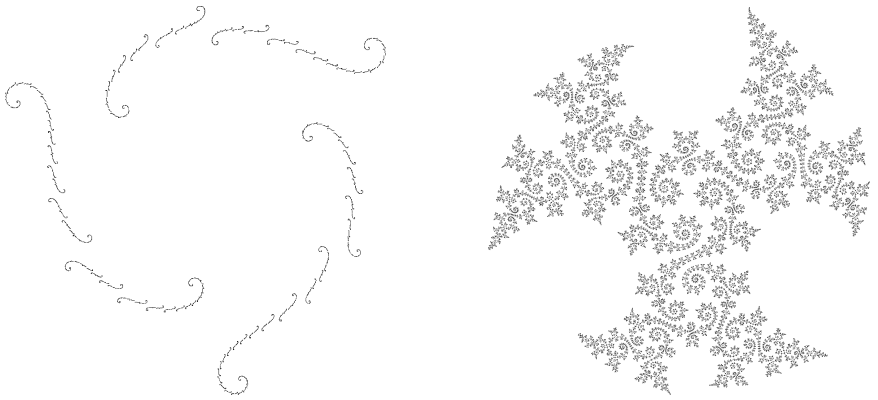

(a) $q = .8 + .3i$, $\delta \approx 1.09$ (b) $q = .8 + .44i$, $\delta \approx 1.29$

Figure 4.2: The limit sets for two of the $\mathbb{Z}_3$ symmetric genus two Schottky groups that arise when investigating $n = 3$ Rényi entropies. The parameter q defining the groups is an eigenvalue of one of the generators as specified in 4.2. We give the value of the Hausdorff dimension δ for these two limit sets, computed using the methods of 4.5.

foray into (Mumford, Series, and Wright, 2002).

Secondly, we note that the size of the terms in the sum 4.30 is controlled primarily by the factor $|\gamma'(w)|^{\Delta}$, which describes how things scale under the action of γ (in the flat or round metric on the Riemann sphere, not the metric pulled back from $\mathcal{B}$). Given some small set near w, the characteristic length of its image under γ is scaled by $|\gamma'(w)|$, and its area is scaled by $|\gamma'(w)|^2$, so it is natural generalise this, and say that $|\gamma'(w)|^{\Delta}$ characterises the scaling in a Δ-dimensional notion of measure, where Δ can be any positive real number. The convergence of the sum is therefore determined by the trade-off between the accumulation of many points at the limit set, and the shrinking of Δ-dimensional measure associated to images at those points. The critical dimension will occur when these two effects precisely balance, which is when the limit set itself can be assigned a Δ-dimensional measure invariant under Γ. Hopefully this discussion makes plausible the following theorem of Patterson (Samuel J Patterson, 1976), Sullivan (Sullivan, 1979; Sullivan et al., 1987) and Bishop-Jones (Bishop and Jones, 1997), the precise statement of which uses the notion of *Hausdorff dimension*, a non-integer dimension defined for fractals in metric spaces.

Theorem 3 (Patterson-Sullivan) *The exponent of convergence δ is equal to the Hausdorff dimension of the limit set Λ of Γ:*

$$\delta = \mathrm{H.dim}(\Lambda) \tag{4.31}$$

This result connecting spectral theory and fractal geometry is certainly beautiful, which would be justification enough to include it in a mathematics paper, but amazingly enough it is also useful. Firstly, it gives us a new tool to intuit how the critical dimension Δ_c depends on the geometry, particularly in certain limits. But more importantly, it provides a method to accurately and efficiently compute δ for any group Γ, which is far better than naïvely solving the bulk Laplace equation numerically, or the method of extracting δ from the asymptotics of the terms in the sums introduced above. We will discuss an algorithm to compute δ in 4.5, and use it to present both numerical and analytic results.

Divergence of the partition function

In this section we will offer one final perspective on the phase transition, to make a direct connection with the CFT argument discussed in section 3. In that section we summed up the contributions to the partition function from of a generalized free field, using the global limit of higher genus blocks with the spectrum. From the bulk point of view, this object is precisely the one-loop partition function of the bulk scalar field ϕ:

$$Z_{\text{GF}} = Z_{\text{1-loop}} = \frac{1}{\sqrt{\det\left(m^2 - \nabla^2\right)}}. \tag{4.32}$$

This makes it apparent that the zero mode should again be visible as a zero eigenvalue of the operator $\nabla^2 - m^2$ (defined with suitable boundary conditions). In this way, the calculations of section 3 put a lower bound on Δ_c.

The bulk computation of this one-loop partition proceeds much as the Green's function computation given above. In particular, one can compute this one-loop determinant using heat kernel methods and a sum over images (Giombi, Maloney, and Yin, 2008). For higher genus surfaces this one-loop determinant can be written as an infinite product

$$Z^{\mathbb{H}^3/\Gamma}_{\text{1-loop}}(\Delta) = \frac{1}{\sqrt{\zeta_\Gamma(\Delta)}}, \quad \text{where } \zeta_\Gamma(\Delta) \equiv \prod_{\gamma\in\mathcal{P}}\prod_{n,\bar{n}=0}^{\infty}\left(1 - q_\gamma^{\Delta+2n}\bar{q}_\gamma^{\Delta+2\bar{n}}\right). \tag{4.33}$$

Here, q_γ is the smaller eigenvalue of γ, as previously introduced. The product is over primitive conjugacy classes $\gamma \in \mathcal{P}$ of the group Γ; these are conjugacy classes of elements which cannot be written as a power γ^n of another element with $n > 1$, and q_γ is the smaller of the eigenvalues of γ when written as an $SL(2,\mathbb{C})$ matrix. Note that this definition counts γ^{-1} separately from γ, so that terms come in matching

pairs. We have written the one-loop partition function terms of the *Selberg zeta function* ζ_Γ associated to the group Γ, as defined in (S. Patterson, 1989)[21]. We note that the product in 4.33 converges in the same right half-plane $\mathrm{Re}\,\Delta > \delta$ as the image sums we have already introduced. In fact, the Selberg zeta function can be analytically continued to an entire function, with zeros precisely at the eigenvalues and resonances of the Laplacian on $\mathcal{M}$, as expected (Samuel J Patterson, Perry, et al., 2001). In particular, the first resonance, corresponding to the phase transition of interest, leads to the one-loop partition function diverging as $(\Delta - \Delta_c)^{-1/2}$. This is the divergence found from the CFT analyses of (Belin, Keller, and Zadeh, 2017) and 4.3.

The product 4.33 has a simple geometric interpretation in terms of the closed geodesics on the bulk manifold $\mathcal{M} = \mathbb{H}^3/\Gamma$. Since Γ is the fundamental group of $\mathcal{M}$, its conjugacy classes are in one-to one correspondence with homotopy classes of closed loops in the bulk, and in a hyperbolic manifold, there is a unique closed geodesic in each class. The primitive conjugacy classes $\mathcal{P}$ correspond to *prime* geodesics that do not trace over their image multiple times. The geometric parameters associated to a closed geodesic are its length ℓ_γ, and its twist θ_γ, the angle through which a normal vector gets rotated after being parallely transported around the curve, and are related to the associated conjugacy class of Γ by $q_\gamma^2 = e^{-\ell_\gamma + i\theta_\gamma}$. The convergence of the product is therefore controlled by the asymptotics of the length spectrum of the bulk manifold. A precise statement of this is given by the prime geodesic theorem, so called because of its close analogy with the prime number theorem (provable using the analytic properties of the Selberg and Riemann zeta functions respectively):

Theorem 4 (Prime geodesic theorem) *The prime geodesic counting function $\pi_{\mathcal{M}}(\ell)$, defined as the number of prime geodesics of length at most ℓ, satisfies the asymptotic formula*

$$\pi_{\mathcal{M}}(\ell) \sim \frac{e^{\delta\ell}}{\delta\ell} \text{ as } \ell \to \infty. \tag{4.34}$$

In this way, the instability of the scalar field is controlled by the asymptotic properties of the spectrum of very long geodesics. This relation between the spectrum of the

[21] There are several closely related definitions of the Selberg zeta function. The definition we have given is appropriate for hyperbolic three-manifolds where Γ is a Kleinian group; another, more common definition is in the context of hyperbolic surfaces, where Γ is a Fuchsian group, so q_γ is real, and the product over $\bar{n}$ is absent.

Laplacian and the lengths of closed geodesics is a special case of the Selberg trace formula (or an appropriate generalization).

Consideration of the partition function leads to an alternative approach to the computation of δ, which we will not pursue further here, by numerically computing the Selberg zeta function, which can be done efficiently (though not directly from the product definition), and locating its zeros.

When Γ is Fuchsian

In the case when the group Γ is Fuchsian, i.e. when all elements are in $SL(2,\mathbb{R})$ and so fix the real line (perhaps after conjugation with some Möbius map, for example fixing the unit circle instead), the discussion simplifies somewhat. Instead of requiring the full three dimensional geometry, all the main results discussed here can be reduced to the two-dimensional slice Σ fixed by complex conjugation. In this section we briefly describe this reduction and its consequences.

Fuchsian groups are, in many circumstances, the most physically interesting cases, primarily because they correspond to geometries that have a real Lorentzian description. Interpreting the action of complex conjugation as a time-reversal symmetry, the slice Σ fixed by time-reversal has vanishing extrinsic curvature, and hence can be interpreted as an initial Cauchy surface for Lorentzian evolution. Very explicitly, the Euclidean bulk $\mathcal{M}$ can be written as

$$ds^2 = d\chi^2 + \cosh^2\chi \; d\Sigma^2 \tag{4.35}$$

where $d\Sigma^2$ is the hyperbolic metric on the $\chi = 0$ slice Σ. The Lorentzian geometry (or, rather, a patch of it) is obtained by analytic continuation $\chi \to it$. This gives a locally AdS_3 solution to the equations of motion in an FRW-like coordinate system, where the spatial slices have constant negative curvature.

It is important to note that while all Fuchsian groups have a reflection symmetry, and corresponding Lorentzian interpretation, the converse is not true: a non-Fuchsian Schottky group may have a time-reflection symmetry and good Lorentzian continuation. To take one example, the pure entangled state on three copies of the CFT obtained by the path integral on a pair of pants is, for certain moduli, dual in the Lorentzian section to disconnected copies of pure AdS and BTZ (Maxfield, Ross, and Way, 2016), but the corresponding (connected) Euclidean geometry is not described by a Fuchsian group.

The bulk metric is not static, so to simplify the Laplacian it is not as straightforward as choosing a time-independent ansatz. But it is not much harder than that; instead, look for a separable eigenfunction $F(\sigma, \chi) = f(\sigma)g(\chi)$, finding that if f is an eigenfunction of the Laplacian on the $\chi = 0$ slice with eigenvalue $\Delta(\Delta - 1)$, and obeys the correct AdS boundary conditions, then

$$F(\sigma, \chi) = (\operatorname{sech} \chi)^{\Delta} f(\sigma) \tag{4.36}$$

is an eigenfunction of the full handlebody Laplacian with eigenvalue $\Delta(\Delta - 2)$, with the correct boundary conditions. From this, the critical dimension of the handlebody is determined by the bottom of the spectrum of the slice Σ, and computing the actual profile of the zero mode is no longer a three-dimensional problem.

4.5 Results for the critical dimension Δ_c

We have seen that a scalar field on a handlebody $\mathbb{H}_3/\Gamma$ will be unstable if the dimension is sufficiently small: $\Delta < \Delta_c$. We now turn to an explicit computation of the critical dimension Δ_c, which will be a function of the moduli. A direct approach, where one studies the Laplacian directly on the geometry of interest, is a complicated numerical task. Our approach will be to instead use 3 to calculate Δ_c from the Hausdorff dimension of the limit set of Γ.

Using this, we will obtain analytic results for Δ_c near the boundary of moduli space, as well as analytic bounds on Δ_c in the interior of moduli space. We will also obtain accurate numerical results. Our main tool will be an algorithm due to McMullen (McMullen, 1998), which we now describe.

McMullen's algorithm

This section is somewhat technical and is not necessary to understand the results described in later sections. Readers who are not interested in the details of how the results are obtained can safely skip to section 4.5.

To begin, will we need to introduce an additional structure on the limit set Λ: a Γ-invariant δ-dimensional measure μ. A measure μ (though we will not give a precise definition here) allows us to integrate functions on the limit set, in particular assigning a number $\mu(E) = \int_E d\mu \geq 0$ to subsets $E \subseteq \Gamma$, providing a measure of the 'content' of E. We require the additional property that it transforms as a

δ-dimensional density under element of the group[22]:

$$\mu(\gamma(E)) = \int_E |\gamma'|^\delta d\mu \quad \text{for } \gamma \in \Gamma \tag{4.37}$$

A nontrivial measure with this property exists when, and only when δ equals the Hausdorff dimension of Λ (in which case it is unique, up to normalization, for Schottky groups). The only feature of the right hand side that we require is that it is bounded by the measure of the set $\mu(E)$, times the extrema of the integrand $|\gamma'|^\delta$:

$$\mu(E) \inf_{w\in E} |\gamma'(w)|^\delta \le \mu(\gamma(E)) \le \mu(E) \sup_{w\in E} |\gamma'(w)|^\delta \tag{4.38}$$

McMullen's algorithm works by splitting the limit set into a finite number of disjoint pieces E_i, and attempting to approximate (or bound) the value of the measure μ on each of these pieces, $\mu_i = \mu(E_i)$. For a detailed explicitly worked example of this and the following, see 4.3. We begin by imposing (4.37). Specifically, suppose we have some Möbius map $\gamma_i \in \Gamma$, and one of the pieces E_i, whose preimage under γ_i is the union of some pieces $E_{j_1}, E_{j_2}, \ldots E_{j_n}$ (one of which may be E_i itself):

$$E_i = \bigcup_{k=1}^{n} \gamma(E_{j_k}) \tag{4.39}$$

If we pick some points $z_j \in E_j$, then (4.37) implies that

$$\mu_i \approx \sum_{k=1}^{n} |\gamma_i'(w_{j_k})|^\delta \mu_{j_k} \,. \tag{4.40}$$

This is not an exact equality because the scale factor $|\gamma_i'|$ is not constant on the limit set. However, by taking the E_i to be small $|\gamma_i'|$ will be approximately constant, so the error will be small. To be more precise, we may replace the factors of $|\gamma_i'(w_j)|$ by upper or lower bounds on this scaling over the set E_j, and use 4.38 to replace the approximate equation by inequalities.

With an appropriately chosen partition $\{E_i\}$ and maps γ_i, a similar argument can be repeated for every i. Our approximate formula 4.40 can then be written in terms of a square matrix T, the *transition matrix*, whose entries T_{ij} are equal to $|\gamma_i'(z_j)|$ for each of the E_j in the preimage $\gamma_i^{-1}(E_i)$, and zero otherwise. Invariance of the measure

[22]Here, $|\gamma'|$ computes the local scaling of lengths under the map γ; we may use any metric on the boundary Riemann sphere for this purpose, for example a round metric, and the result for δ is insensitive to this choice. For practical computations the flat metric is often most convenient, in which case $|\gamma'|$ is the absolute value of the derivative of the Möbius map. It is then simplest to require that the point at infinity is not in Λ, which we will implicitly assume.

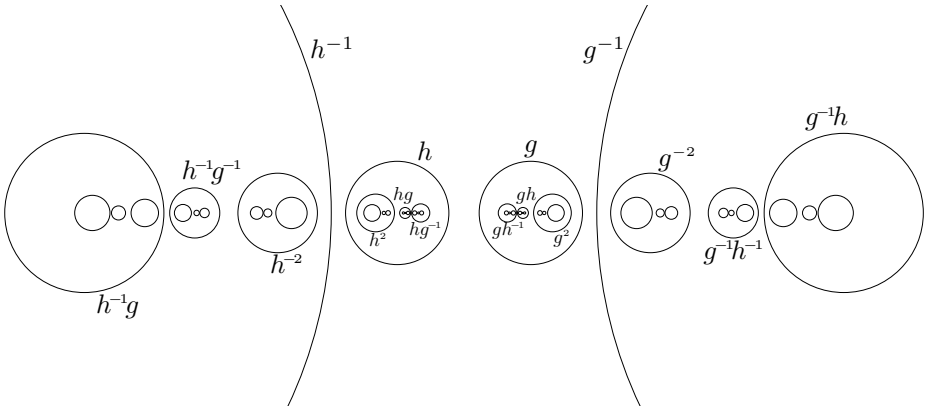

Figure 4.3: An example of computing the transition matrix for McMullen's algorithm, in the case of a Kleinian group freely generated by two loxodromic elements g, h, so that $\mathbb{C}^*/\Gamma$ is a genus two surface. In the figure, we have drawn a fundamental domain for Γ, the exterior of the four outermost circles (those corresponding to g^{-1}, h^{-1} are not shown in their entirety). Break the limit set into the four pieces E_γ contained in each of these circles, labelled by $\gamma = g, h, g^{-1}, h^{-1}$ corresponding to the element of the group that maps the fundamental domain to the interior of the circle, and choose points $w_\gamma \in E_\gamma$, for example the attractive fixed point of γ. The piece of the limit set E_g can be broken up into three disjoint pieces, inside the circles labelled g^2, gh and gh^{-1}, which are the images under g of E_g, E_h and $E_{h^{-1}}$ respectively. The scalings of these limit sets under the action of g go into the top row of the transition matrix:

$$T = \begin{pmatrix} |g'(w_g)| & |g'(w_h)| & 0 & |g'(w_{h^{-1}})| \\ |h'(w_g)| & |h'(w_h)| & |h'(w_{g^{-1}})| & 0 \\ 0 & |(g^{-1})'(w_h)| & |(g^{-1})'(w_{g^{-1}})| & |(g^{-1})'(w_{h^{-1}})| \\ |(h^{-1})'(w_g)| & 0 & |(h^{-1})'(w_{g^{-1}})| & |(h^{-1})'(w_{h^{-1}})| \end{pmatrix}$$

The other three rows repeat the same exercise for the other three regions, and finding δ such that the spectral radius of T^δ is unity gives an approximation for the Hausdorff dimension. This can be refined by breaking the limit set up into the $3 \times 4^{n-1}$ regions E_γ labelled by words of length n in g, h, g^{-1}, h^{-1}, and applying the δ-invariance imposed by considering the preimage of E_γ under the first element (g, h, g^{-1}, or h^{-1}) appearing in the word γ. Then T will be a sparse matrix, with three nonzero elements in each row and column, and the algorithm has error decreasing exponentially with n. The figure includes labels for words of length two, but also shows the images of circles under words of length three (unlabelled).

is then the statement that μ_i is a unit eigenvector of T^δ, where the power is taken element-wise. Since the matrix T^δ has nonnegative entries it is guaranteed to have a unique eigenvector with positive components; furthermore this is the eigenvector with largest eigenvalue.[23] If μ is to satisfy the invariance criterion, this eigenvalue should be one. Thus to find the Hausdorff dimension, we find the value of δ such that the largest eigenvalue of T^δ equals one.[24] An analogous result holds when we replace the approximate equations by inequalities, so by choosing upper or lower bounds on the transition matrix elements, we can obtain rigorous bounds on δ.

For a given partition $\{E_i\}$ of the limit set this gives an estimate for δ, and to obtain a more accurate estimate we can refine the partition into a larger number of pieces. With an appropriate refinement, the result converges rapidly to the Hausdorff dimension, and in practice it is sufficient to use a rather coarse partition of Λ. Although this discussion is rather abstract, the explicit implementation of this algorithm is quite straightforward; see figure 4.3 for a simple example.

Analytic results

Our first analytic results are for Fuchsian groups; this includes the surfaces described in 4.2 with real cross-ratio $0 < x < 1$. In this case we note that the limit points must all lie on the real axis of the w-plane, corresponding to the slice fixed by time-reflection symmetry. Since Λ is a subset of the one-dimensional line, it must have dimension $\delta \leq 1$. Thus $\Delta_c \leq 1$. So the only potentially unstable fields are those with $\Delta < 1$, which correspond to bulk scalars which are quantized using alternate boundary conditions.

In the rest of this section we will focus on the case of the genus $g = n - 1$ surfaces with $\mathbb{Z}_n$ symmetry, described in 4.2, relevant for computing the nth Rényi entropy of a pair of intervals. We begin by applying the above algorithm in this case at the coarsest level of approximation, to obtain analytic bounds. These are especially useful at the edge of moduli space where $x \to 0$, because in this limit the different pieces of the limit set become well-separated and small, so the scale factor does not vary much over it. These bounds thus become tighter and tighter as $x \to 0$.

[23] This is the Perron-Frobenius theorem. Since some of the entries of T are zero, we must also require that T is irreducible. This means, roughly speaking, that when we apply T repeatedly all of the regions E_i will eventually mix.

[24] This δ exists and is unique since the spectral radius of T^δ decreases monotonically as δ increases from zero to infinity.

As described in 4.2, we can extend the group Γ to $\hat{\Gamma}$, generated by $R : w \mapsto e^{2\pi i/n} w$ and $S : w \mapsto \frac{w-\zeta}{\zeta w - 1}$, by including a dihedral group of holomorphic automorphisms. This extension of the group does not alter the limit set, and a Γ-invariant measure constructed on it will also be invariant under $\hat{\Gamma}$. Using this, we will divide the limit set into n pieces, all related by the $\mathbb{Z}_n$ symmetry R, and hence having equal measure, and use the mapping under S to constrain the dimension of this measure. This is somewhat simpler than using the original presentation of Γ such as in 4.3.

There are n pieces of the limit set E_k, each centred at a root of unity $e^{2\pi i k/n}$, with size of order θ^2 for small θ, and related to each other by R. A simple way to show this is by constructing a fundamental domain for $\hat{\Gamma}$, bounded by the radial lines from the origin at angles $\pm\pi i/n$, related by action of R, and a circle C_n mapped to itself by S, centred at $\sec\theta$ with radius $\tan\theta$ (recall $\zeta = \cos\theta$). Then E_n is the part of the limit set inside this circle, and the remaining E_k are inside corresponding circles $C_k = R^k(C_n)$ obtained by rotations by angle $2\pi k/n$. This immediately bounds the size of the limit set by the size of the circles C_k, of order θ, which is a sufficiently strong result for our immediate purposes[25].

By the $\mathbb{Z}_n$ symmetry, and the fact that the action of R doesn't scale ($|R'| = 1$), the sets E_k must all have equal measure $\mu_k = \mu$. We can then apply the action of S, which maps $E_1, E_2, \cdots, E_{n-1}$ onto E_n. The amount by which E_k scales under S can be computed from

$$S'(e^{2\pi i k/n} + O(\theta^2)) = \frac{\theta^2}{(e^{2\pi i k/n} - 1)^2} + O(\theta^4) \tag{4.41}$$

so 4.40, equating the sum of the scaled measures of $E_1, E_2, \cdots, E_{n-1}$ to the measure of E_n, gives us

$$\sum_{k=1}^{n-1} \left(\frac{\theta}{2\sin\frac{\pi k}{n}} \right)^{2\delta} = 1 + O(\theta^2). \tag{4.42}$$

This requires δ to tend to zero as θ goes to zero, and solving to leading order for small δ, we get the result that

$$\delta \sim \frac{\log(n-1)}{2\log|\theta|^{-1}} + O\left(\frac{1}{(\log|\theta|)^2}\right) \quad \text{as } \theta \to 0 \tag{4.43}$$

with the first corrections coming from solving the equation to higher order in δ, rather than the order θ^2 variation of the scaling relation giving the correction to

[25] To obtain an improved bound, note that E_n must be contained in a smaller set, namely the union of the interiors of $S(C_k)$, which is concentrated in a region with size of order θ^2.

4.42. With only minor modifications, this derivation continues to apply if we allow θ to be complex, which is why we have included the modulus in the result. The case $n = 3$, with real θ, was treated in (McMullen, 1998), though instead of using S, that paper uses a reflection, which requires θ to be real[26].

To facilitate comparison between different values of n, we write this in terms of the cross-ratio x. With the monodromy methods outlined in 4.2, the map from x to the Schottky parameter θ can be computed as a series expansion, with the leading order result that $\theta = \frac{\sqrt{x}}{n}(1 + O(x))$:

$$\delta = \frac{\log(n-1)}{\log|x|^{-1}} + O\left(\frac{1}{(\log|x|)^2}\right) \quad \text{as } x \to 0 \tag{4.44}$$

Note that this is the asymptotic behaviour for fixed n as $x \to 0$, but clearly must break down if n is parametrically large. In the first instance, it is not self-consistent if nx is of order one, since in that case the leading order term in the expansion would be of order one. But we can take a different order of limits to understand what happens at large n.

Starting with the scaling relation 4.42 at small fixed x, and naïvely taking the large n limit term by term, we arrive at

$$2\sum_{k=1}^{\infty}\left(\frac{\sqrt{x}}{2\pi k}\right)^{2\delta} \approx 1 \tag{4.45}$$

where the factor of two is to count contributions both from fixed k and fixed $n - k$. The first thing to notice is that the tail of the series decays as $k^{-2\delta}$, so convergence of the sum immediately requires $\delta > \frac{1}{2}$. Summing the series, we arrive at

$$\frac{1}{2}\left(\frac{4\pi^2}{x}\right)^{\delta} \approx \zeta(2\delta) \tag{4.46}$$

from which we see that the left hand side is large for small x, since δ cannot be small. The zeta function must therefore be close to the pole at $\delta = \frac{1}{2}$, and we can find a perturbative solution:

$$\lim_{n\to\infty}\delta = \frac{1}{2} + \frac{\sqrt{|x|}}{2\pi} + O(x) \tag{4.47}$$

This result should be interpreted as the limit of δ as $n \to \infty$, for fixed but small x. It turns out this naïve argument is, in essence, correct, and can be made completely

[26] For comparison, the parameter θ used in that paper is half of the θ used here.

precise by repeating the argument for the group generated by one parabolic element and one elliptic element of order 2, equivalent to Theorem 3.6 of (McMullen, 1998) (for real cross-ratio). In fact, the result that the Hausdorff dimension does not go to zero as $x \to 0$ is a consequence of a general result, that $\delta > \frac{1}{2}$ whenever the group in question contains a parabolic element (Corollary 2.2 in (McMullen, 1999)).

This Schottky group with R parabolic, instead of elliptic order n, corresponds to the '$n = \infty$' version of the geometry described in 4.2. This complex one-dimensional family of Kleinian groups is known in the mathematics literature as the 'Riley slice' of Schottky space. It is a little tricky to think about the $n \to \infty$ limit of the handlebody, bounded by a Riemann surface of infinite genus, but it is rather simpler to understand the geometry after taking a quotient by the $\mathbb{Z}_n$ replica symmetry, as suggested in (Lewkowycz and J. Maldacena, 2013). The boundary of this geometry is just the original Riemann sphere, and the bulk has conical defects, of opening angle $2\pi/n$, going from 0 to x and from 1 and ∞. Taking a formal analytic continuation of the geometry to $n = 1$, the conical defects become the Ryu-Takayanagi surface, lying on geodesics (Lewkowycz and J. Maldacena, 2013), but we are taking the opposite limit, in which the defects become cusps, in particular receding to infinite proper distance.

At this point, let us pause briefly to understand the physical consequences. This result means that the $x \to 0$ and $n \to \infty$ limits of the critical dimension do not commute, so that while for any fixed n, any dimension of scalar will be stable for sufficiently small x, if there is a scalar of dimension less than $1/2$, it will be subject to the phase transition for *any* cross ratio, if n is taken sufficiently large. Note that this is all in a limit where we have taken c to infinity first, and new behaviour dominated by quantum corrections may take over when n is parametrically large in c. In particular, the large n limit of the Rényi entropy is controlled by the largest eigenvalue of the reduced density matrix, or the 'ground state energy' of the modular Hamiltonian $H_A = -\log \rho_A$, but it is unclear whether the semiclassical description is sensitive to a single lowest eigenvalue, or a dense collection of parametrically many low-lying eigenvalues of H_A.

To conclude the discussion of analytic results, let us briefly describe the other limit of the geometry, when $x \to 1$, corresponding to the horizon sizes of the multi-boundary wormhole becoming small. This limit is of less direct physical interest, since it is well past the first-order phase transition of the partition function, so is not the dominant saddle-point geometry, but is nonetheless useful to hone intuition. In

the case that x is real, so Γ is a Fuchsian group, it is sufficient to describe the bulk geometry in terms of the two-dimensional hyperbolic surface making up the $t = 0$ slice. In this limit, it is helpful to separate the geometry into the exterior pieces, lying between each boundary and a horizon, and the 'convex core' or 'causal shadow' region linking them together, bounded by the n horizons. In the $x \to 1$ limit, the horizons become very small, and the centre of the geometry recedes down a long, narrowing neck. The core then approximates the geometry of the negatively curved metric on some compact surface with punctures, though the punctures do not quite pinch off, rather reaching a minimum radius at the narrow horizons where they join to the exterior funnels. On such a surface, the critical dimension approaches one, the maximal eigenvalue of the Laplacian (on the $t = 0$ slice, not the Laplacian in the three-dimensional bulk: see 4.4) being small and positive. The corresponding eigenfunction is roughly constant on the convex core, and small in the exterior funnels, with an interpolation over the long, narrow necks connecting them. The asymptotic behaviour of Δ_c in this limit can be computed by directly approximating this Laplace eigenfunction (the zero mode of the instability) (Dodziuk et al., 1987).

Taking $x \to 1$ through complex values is much more complicated, so we will not be able to say much about it. Schottky space, perhaps parameterized in this case by the values of ζ corresponding to some complex x, itself has a complicated fractal boundary, and the features of the handlebody depend sensitively on how this boundary is approached. This is a deep and beautiful subject, but goes far beyond the scope of this article. In any case, the numerical computations we describe next show that it is possible to obtain dimensions $\delta > 1$ in this limit, for example the limit sets illustrated in 4.2.

Numerical results

These analytic results are very useful to understand the behaviour of the critical dimension at the edges of moduli space, and as genus is varied, but McMullen's algorithm is also useful to quickly compute the Hausdorff dimension numerically, to many digits of precision. We conclude by presenting the results of these computations.

Firstly, 4.4 plots the Hausdorff dimension as a function of cross-ratio x for various values of n, including also the limit as $n \to \infty$. With this parameterization, for generic values of x, the convergence of the algorithm is remarkably rapid. Indeed, the plot includes shaded regions to indicate the rigorous bounds obtained by applying

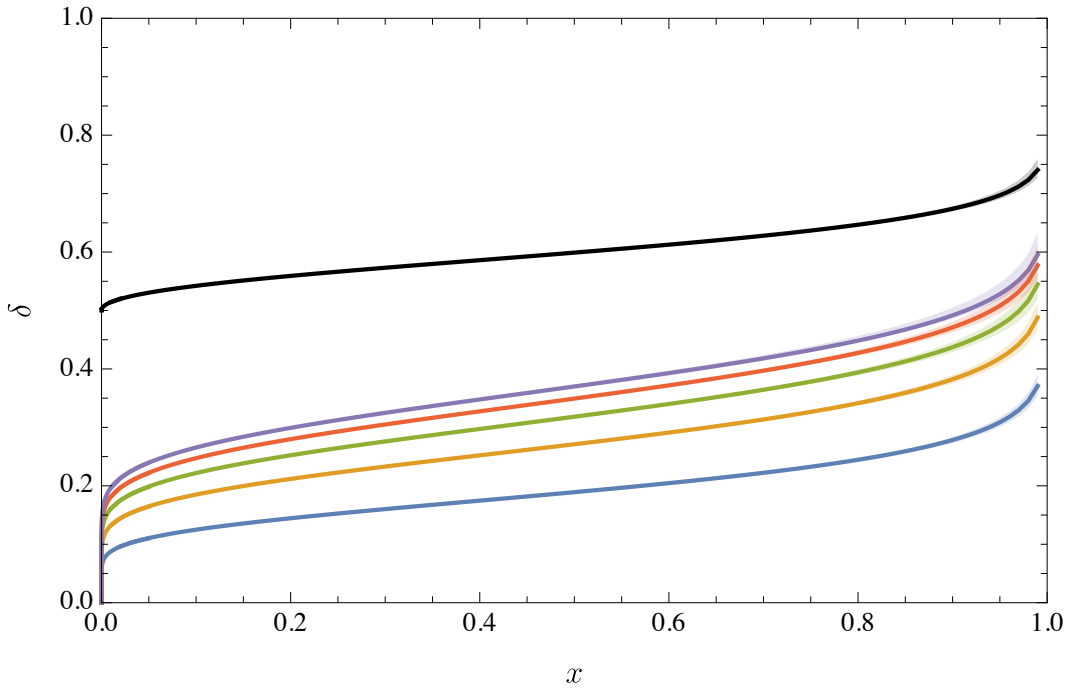

Figure 4.4: The critical dimension $\Delta_c = \delta$ as a function of cross-ratio x for the handlebodies corresponding to the Rényi entropies of a pair of intervals. From top to bottom, the curves correspond to genus 2, 3, 4, 5, 6, and finally the $n \to \infty$ result in black. The shading visible on the right side of the plot indicates the bounds achieved by applying McMullen's algorithm at the crudest level of approximation.

the algorithm at the crudest level. These are computed by numerically solving the equation 4.42, but replacing the terms on the left hand side with upper or lower bounds for $|S'(w)|^\delta$ over $w \in E_k$, rather than the estimates used there. The allowed regions for δ are in many cases not even visible until x is rather close to 1. Refining further, the algorithm gives results with ten or more digits of precision in a fraction of a second on a laptop of modest specifications. In fact, by far the larger source of computing time and error comes from the conversion between the cross-ratio and Schottky variables, rather than the algorithm to compute δ from the group generators.

A physically motivated value to consider is at the boundary with the Hawking-Page phase transition $x = \frac{1}{2}$, which will give the maximum value of Δ_c within this class of geometries, while in the dominant phase. At genus 2 ($n = 3$), this value is $\Delta_c = 0.189124003$, which is rather close to (and the correct side of) the bound $\Delta_c \geq 0.18912109$ obtained in 4.3 from refining the CFT methods of (Belin, Keller, and Zadeh, 2017). Staying at $x = \frac{1}{2}$ and increasing the genus, we find that the

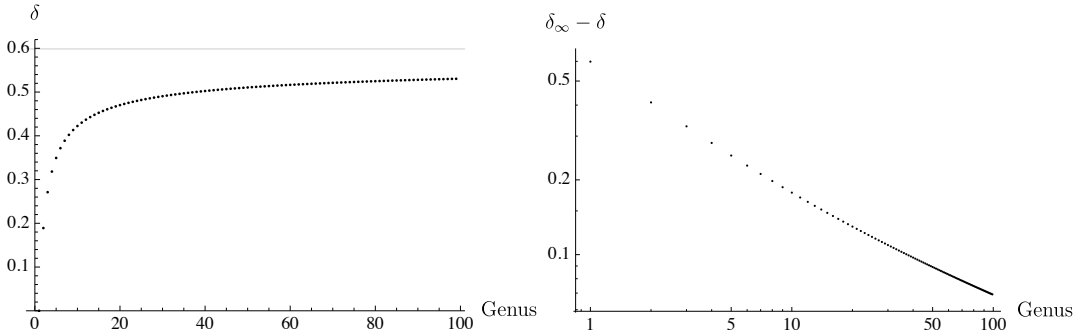

Figure 4.5: The critical dimension for the $x = \frac{1}{2}$ Rényi surface as a function of replica number n. The asymptote is the computed limit as $n \to \infty$. On the right is a log-log plot showing convergence to this value.

critical dimension increases rapidly at first, before slowly approaching the limiting value $\Delta_c \to 0.599$ as $n \to \infty$, as shown in 4.5.

In our last plot, 4.6, we indicate how the Hausdorff dimension behaves for complex values of the cross-ratio, for genus two. Note that this is invariant under inversion in the circle of unit radius centred at one. This is because the extended groups $\hat{\Gamma}$ corresponding to these geometries are the same, though Γ consists of different subgroups in each case. From the geometric point of view, taking the $\mathbb{Z}_n$ quotient of the handlebody gives the same geometry, with conical defects at the fixed points, though the original geometries are distinct (being branched around the defects in different ways). This relates a cross-ratio $0 < x < 1$ with a negative cross-ratio $-\frac{x}{1-x}$, which corresponds to swapping the location of twist and anti-twist operators, relevant for computing Rényi negativity of two disjoint intervals (Calabrese, J. Cardy, and Tonni, 2012). The correspondence between the geometries implies a correspondence between the classical limits of Rényi entropy and Rényi negativity for two intervals, and taking an analytic continuation to $n = 1$, the logarithmic negativity of two intervals must vanish (to leading order in c) in the regime $x < \frac{1}{2}$ where this geometry dominates the path integral.

Finally, it is interesting to ask what the largest possible value of Δ_c could be for a geometry that dominates the path integral. A lower bound (precluding surprising new symmetry-breaking phases that dominate the path integral) comes from our numerics for the infinite genus limit, giving examples where Δ_c as large as .599 can be achieved. An interesting result that may bound this in the other direction comes from (Hou, 2016), showing in particular that every Riemann surface admits a uniformization by a Schottky group of Hausdorff dimension less than one. As a

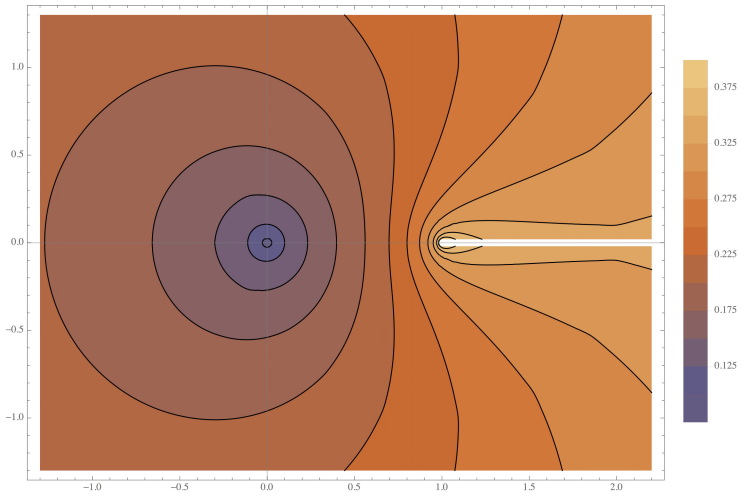

Figure 4.6: The Hausdorff dimension of the $\mathbb{Z}_3$ symmetric genus two handlebody, as a function of the (complex) cross-ratio x. Note that the Hausdorff dimension goes to zero at the origin ($x = 0$) and approaches one as $x \longrightarrow 1^-$ along the real axis.

heuristic, matching our expectations in limits of moduli space, the dominant saddle-point seems to be that with minimal Hausdorff dimension, so this is suggestive, though not conclusive, that there may never be a dominant geometry with $\Delta_c > 1$.

Chapter 5

Going Further in AdS_3

This chapter describes some of our results in AdS_3 beyond what's in [DMMM], specifically related to analytic continuation and then to classifying multiboundary wormholes.

5.1 Analytic Continuation in AdS_3

This section rapidly surveys Kiril Krasnov's program for analytic continuation in AdS_3 and then comments on how to complete it. Recall that the isometry groups of AdS_3 have the following forms

$$\text{Isom}\left(AdS_3^L\right) = \frac{SL(2,\mathbb{R}) \times SL(2,\mathbb{R})}{\mathbb{Z}_2} \tag{5.1}$$

$$\text{Isom}\left(AdS_3^E\right) = PSL(2,\mathbb{C}). \tag{5.2}$$

On the highest level, we want a way to go back and forth from Lorentzian quotients to Euclidean quotients. For example, given a subgroup $\Gamma_L \times \Gamma_R \subset SL(2,\mathbb{R}) \times SL(2,\mathbb{R})$ we want a procedure to obtain a subgroup Γ_E of $PSL(2,\mathbb{C})$. And also to be able to go in the reverse direction[1].

It's obvious how to perform analytic continuation when there's a global timelike Killing vector. You just do the normal $t \mapsto -i\tau$. This shows how to analytically continue AdS_3^L to $AdS_3^E \cong \mathbb{H}^3$ and vice-versa. However, it's not obvious how to proceed when there's no global timelike Killing vector, which is the case with non-trivial multiboundary wormholes, where timelike Killing vectors can only be defined

[1]Note that we only need to be able to do this for physically realistic spacetimes. For example, on the Lorentzian side of things, we need Γ_L and Γ_R to have the same number of generators. We also need each of their individual circle pairings to not intersect. Finally we need the two groups to generate a spacetime where the angular momentum of the horizons is below extremal bounds.

for each asymptotic region; not globally. Kiril Krasnov introduced a program to analytically continue multiboundary wormholes in the series of papers (Krasnov, 2000; Krasnov, 2002; Krasnov, 2003).

The first paper in this series showed how to perform analytic continuation when there's a time-reflection symmetry. Recall that time-reflection symmetric MBWs correspond to when $\Gamma_L = \Gamma_R$. On the Lorentzian side of things, his approach follows from the fact that there's a natural embedding of $PSL(2,\mathbb{R})$ into $PSL(2,\mathbb{C})$. On the Euclidean side of things, time reflection symmetry means that $\Gamma_E \subset PSL(2,\mathbb{R}) \subset PSL(2,\mathbb{C})$. Using a diagramatic notation to organize this, with time-reflection symmetric spacetimes the analytic continuation procedure works simply by setting $G = \Gamma_L = \Gamma_R = \Gamma_E$ and identifying

$$AdS_3^L/G \quad \longleftrightarrow \quad AdS_3^E/G. \tag{5.3}$$

It's less obvious how to proceed when the time reflection symmetry is broken. Krasnov argued in (Krasnov, 2002; Krasnov, 2003) that one can go from Euclidean quotients to Lorentzian quotients by performing a Fenchel-Nielsen deformation.

In order to complete the program, one needs to understand how to analytically continue from rotating Lorentzian spacetimes to Euclidean spacetimes. Specifically, this means that one starts with $\Gamma_L, \Gamma_R \subset PSL(2,\mathbb{R})$ and wants to obtain a $\Gamma_E \subset PSL(2,\mathbb{C})$.

There's a natural way to do this using Bers' *simultaneous uniformization theorem* (Bers, 1960). The full details of this theorem are extremely beautiful but outside the scope of this document. For intuition, we recommend reviewing Brock's note (Brock, 2007). In particular, the proof relies on techniques from the theory of quasiconformal mappings. However, for a summary version of the theorem, it says that given closed Riemann surfaces, $\Sigma_1, \Sigma_2 \subset$ one can obtain a Quasi-Fuchsian group $\Gamma_{QF} \subset PSL(2,\mathbb{C})$ so that the quotient $\hat{\mathbb{C}}/\Gamma_{QF}$ gives a disjoint union of the two initial Riemann surfaces. Bers' procedure is unique up to an identification of the fundamental groups $\pi_1(\Sigma_1)$ and $\pi_1(\Sigma_2)$. See figure 5.1 for intuition.

Start with initial closed Riemann surfaces given by the Schottky doubles of Γ_L and Γ_R, given as $\hat{\mathbb{C}}^2/\Gamma_L$ and $\hat{\mathbb{C}}^2/\Gamma_R$. Taking the trivial identification between the fundamental groups of $\hat{\mathbb{C}}^2/\Gamma_L$ and $\hat{\mathbb{C}}^2/\Gamma_R$, the simultaneous uniformization theorem gives a unique way to analytically continue from Lorentzian quotients to Euclidean quotients by identifying

$$AdS_3^L/(\Gamma_L, \Gamma_R) \quad \longmapsto \quad AdS_3^E/\Gamma_{QF}. \tag{5.4}$$

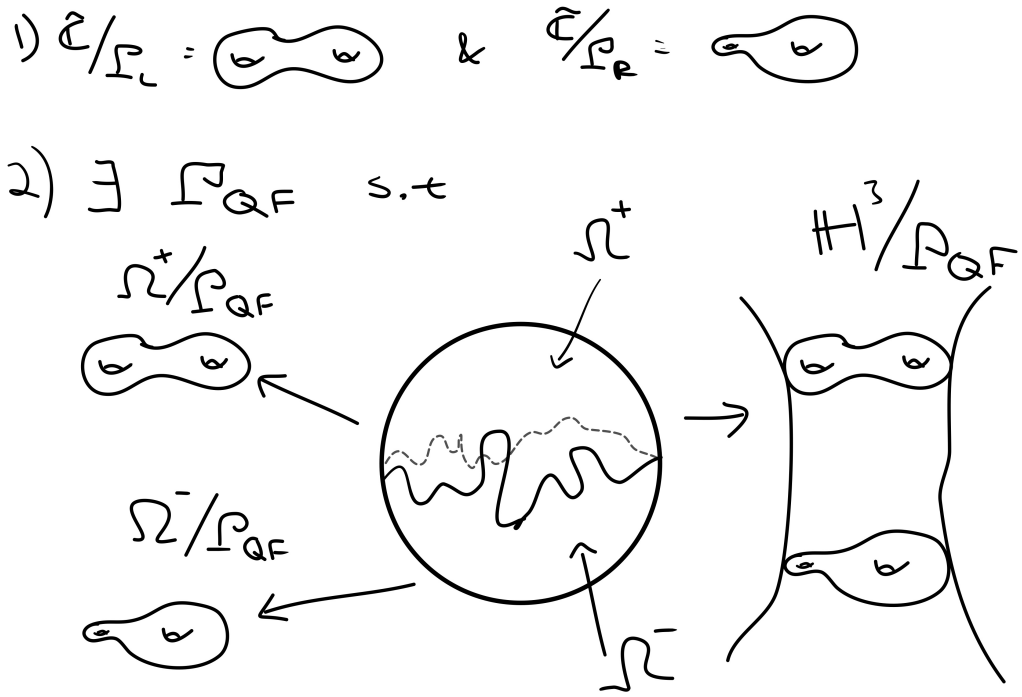

Figure 5.1: This figure sketches how the simultaneous uniformization procedure works. Start with two closed Riemann surfaces. The theorem says there exists a unique $\Gamma_{QF} \subset PSL(2,\mathbb{C})$ and a partition of the complex sphere $\hat{\mathbb{C}} = \Omega^+ \cup \Omega^- \cup \gamma_{QF}$ (γ_{QF} is just a Jordan curve that forms the boundary between Ω^+ and Ω^-) such that the quotients of the Ωs by Γ_{QF} give the appropriate initial Riemann surfaces. The quotient $\mathbb{H}^3/\Gamma_{QF}$ gives a hyperbolic three-manifold that has the topological structure $\mathbb{R}\times\hat{CC}/\Gamma_L$ but limits to Ω^+/Γ_{QF} at positive infinity and to Ω^-/Γ_{QF} at minus infinity.

Note that there are a number of subtleties here that we haven't fully worked out. Especially that here we're applying the simultaneous uniformization theorem to the Schottky doubles rather than to the non-compact surfaces $\mathbb{H}^2/\Gamma_L$ and $\mathbb{H}^2/\Gamma_R$, respectively. In order to obtain a result that matches what we'd expect for the time-reflection symmetric case one needs to solve a highly non-trivial inverse problem. This quick sketch is only an outline for how to complete Krasnov's program.

5.2 Classification of Multiboundary Wormholes (MBWs)

This chapter attempts to classify multiboundary wormholes. Much of this is scattered throughout the literature but there's no resource that puts it all together. First we'll describe the allowed topologies. We saw in chapter 2 that if one starts with a Fuchsian $\Gamma \subset PSL(2,\mathbb{R})$ then the associated Lorentzian spacetime is the quotient

AdS_3^L/Γ which has topology isomorphic to $\mathbb{R} \times \mathbb{H}^2/\Gamma$. We can classify all possible MBW topologies by restricting to Fuchsian groups because going to more general quotients simply turns on angular momentum in the asymptotic regions (reference Krasnov.)

Let's define G to be the genus of the Schottky double $\hat{\mathbb{C}}/\Gamma$, h to be the number of holes in $\mathbb{H}^2/\Gamma$ and b to be the number of boundary regions in $\mathbb{H}^2/\Gamma$. The Schottky double must obey the following formula

$$G = 2h + b - 1. \tag{5.5}$$

Imagine glueing together two copies of a Riemann surface with h holes and b boundaries (disks removed.) Each hole in the single copy leads to two holes in the double, but for every boundary region, the first just makes the surfaces connected and then for each after we get another hole in the Schottky double. The allowed topologies correspond to solutions of $G = 2h+b-1$ where $G, h, b \in \mathbb{N}$. When $b = 0$ the solutions correspond to vacuum diagrams and it's not exactly clear how they'll contribute to the partition function (or if they should.) Table 5.1 demonstrates the pattern.

There are a few things to clarify from this table. The entries with asterisks next to them correspond to vacuum contributions. It's not clear whether or not they should be counted in the partition function. However, it's worth mentioning that strictly speaking they are genuine solutions of Einstein's equations except at a finite number of points that correspond to singularities. Another subtlety with these solutions is that they don't obey the equation $G = 2h + b - 1$. They don't have any boundaries so they obey the formula $G = 2h$. Also note that their Schottky doubles will be disconnected. These vacuum contributions correspond to Fuchsian Schottky groups where the paired circles are all kissing. See Figure 5.3 for an example.

Also note that this argument only described the topology. There are still interesting questions regarding the geometry, such as finding constraints on the angular momentum of different horizons (similar to the extremal Kerr and BTZ black holes.) This latter point will be bounded geometrically by entanglement inequalities via the Ryu-Takayanagi correspondence. There's also a question about the placement of horizons (both inner and outer.) Finally, note that this table only describes the multiboundary wormhole solutions. There are other possible solutions, such as the thermal solutions which strictly speaking aren't always multiboundary wormholes.

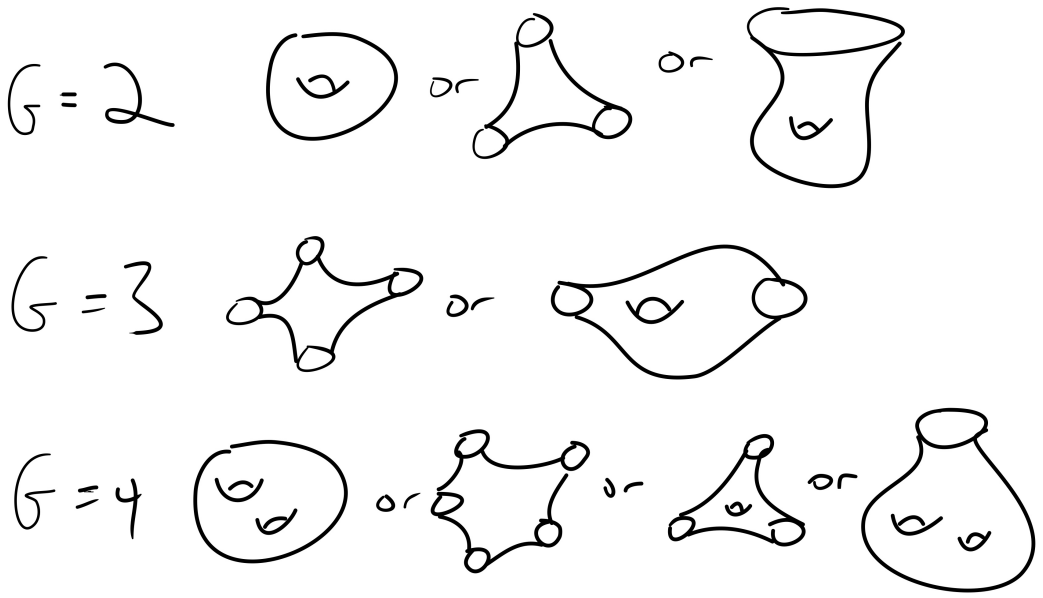

Figure 5.2: This figure shows the options for the first few entries in table 5.1. For $G = 2$ there's a vacuum contribution given as a closed torus, the three boundary wormhole and the torus wormhole. Each of these surfaces should be thought of as the $t = 0$ slice with the global (Lorentzian) topology thought of as a product of each of one of these surfaces times $\mathbb{R}$.

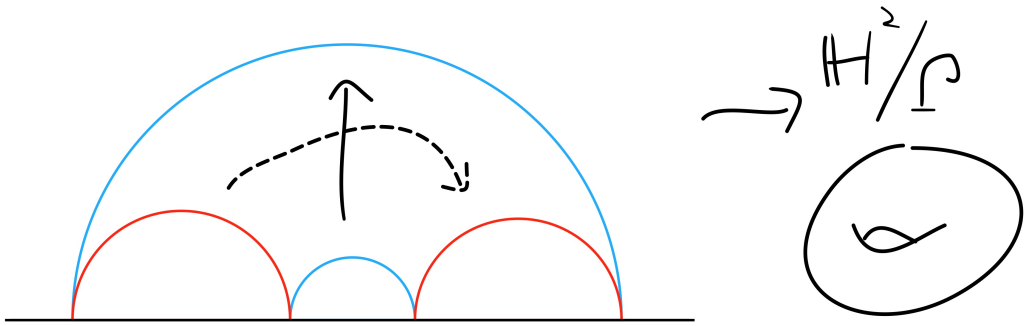

Figure 5.3: This shows the torus vacuum contribution, showing the topology obtained when all of the Schottky circles kiss. The result is a closed Riemann surface with the topology of a torus. In this figure, the blue circles and red circles are identified, respectively.

G	h	b	Name of MBW
0	0	0	Pure AdS_3
1	0	2	BTZ black hole
*2	1	0	Torus vacuum contribution
2	0	3	Three boundary wormhole
2	1	1	Torus wormhole
3	0	4	Four boundary wormhole
3	1	2	
*4	2	0	2-torus vacuum contribution
4	0	5	Five boundary wormhole
4	1	3	
4	2	1	
5	0	6	Six boundary wormhole
5	1	4	
5	2	2	
*6	3	0	3-torus vacuum contribution
6	0	7	Seven boundary wormhole
6	1	5	
6	2	3	
6	3	1	
7	0	8	Eight boundary wormhole
7	1	6	
7	2	4	
7	3	2	
*8	4	0	4-torus vacuum contribution
8	0	9	Nine boundary wormhole
⋮	⋮	⋮	

Table 5.1: This table shows the allowed global topologies of multiboundary wormholes. G is the genus of the Schottky double, h is the number of holes in the $t = 0$ slice and b is the number of boundaries in the $t = 0$ slice. The asterisks correspond to pathological entries with no boundary regions; it's not clear how to treat these in the partition function. Not all entries are named.

Two torus wormhole example

I'd like to describe the power of this classification process through an example. Say one wanted to construct a solution that has one asymptotic region and a two-torus behind the horizon. I've never seen a solution like this in the literature and it would be hard to construct naively. However, with this classification machinery it's easy. In order to have one asymptotic region and two holes we know that $b = 1$ and $h = 2$ so the number of generators we need 4 generators. We now need to choose a circle pairing that gives the right topology. Chasing circle pairings directly gets

unwieldy quickly: even following the identifications with the torus wormhole can be error-prone. To proceed with the two torus wormhole example, we can guess that the circle pairing in figure 5.4 has the right topology.

All that we need to do to check that this gives the right topology is to confirm that the pairing in 5.4 only has one boundary region. Because $G = 4$, the formula $G = 2h + b - 1$ then verifies that there's a two-torus living behind the horizon. In order to check that there's only one boundary region, we need to make sure all of the identifications lead to a single closed cycle. More specifically, the chosen circle pairings tell us that the following points are identified

$$\begin{aligned} a_1 &\sim e_2 \\ a_2 &\sim e_1 \\ b_1 &\sim f_2 \\ b_2 &\sim f_1 \\ c_1 &\sim g_2 \\ c_2 &\sim g_1 \\ d_1 &\sim h_2 \\ d_2 &\sim h_1 . \end{aligned}$$

Furthermore, the boundary segments are (a_2, b_1), (b_2, c_1), (c_2, d_1), etc. If we also write these as identifications we want to make sure there's only one closed cycle starting from a_1 and ending back at a_1. The 'identifications' from the boundary segments are

$$\begin{aligned} a_2 &\sim b_1 \\ b_2 &\sim c_1 \\ c_2 &\sim d_1 \\ d_2 &\sim e_1 \\ e_2 &\sim f_1 \\ f_2 &\sim g_1 \\ g_2 &\sim h_1 \\ h_2 &\sim a_1 \\ &. \end{aligned}$$

If we follow the path that comes from alternating between identifications of points

and identifications of boundary segments we see that there's only one cycle

$$a_1 \sim e_2 \sim f_1 \sim b_2 \sim c_1 \sim g_2 \sim h_1 \sim d_2 \sim e_1 \tag{5.6}$$

$$\sim a_2 \sim b_1 \sim f_2 \sim g_1 \sim c_2 \sim d_1 \sim h_2 \sim a_1. \tag{5.7}$$

This was only a rough sketch but it shows that there is only one boundary region.

One may ask how we guessed at this initial circle pairing? This comes with experience but there are some tricks that make it obvious. For example, for every pair of neighboring circles that are identified there will be an asymptotic boundary. See figure 5.5 for a sketch. No matter how the four black circles are paired there will be at least two asymptotic boundaries here, one that corresponds to the back line segment in between the blue endpoints and another for the other black boundary segment between the green endpoints. In this example, there are 4 generators so the equation $G = 2h + b - 1$ becomes $5 = 2h + b$ which implies that there are an odd number of boundary segments. If one pairs the black circles to their neighbors then we'll be left with the 5-boundary wormhole. If one pairs to opposite black circles then we'll be left with an MBW with three boundary regions and one hole.

Once we know the order of circle pairings to give a desired topology it's easy to find explicit generators using the circle pairing trick from section 3.1. The horizon is given by the word $ABCDA^{-1}B^{-1}C^{-1}D^{-1}$. We can use formula 2.21 to calculate its length ℓ_H

$$\ell_H = 2\cosh^{-1}\left(\frac{\mathrm{Tr}(ABCDA^{-1}B^{-1}C^{-1}D^{-1})}{2}\right). \tag{5.8}$$

This example was meant to demonstrate how efficiently the MBW machinery fits together.

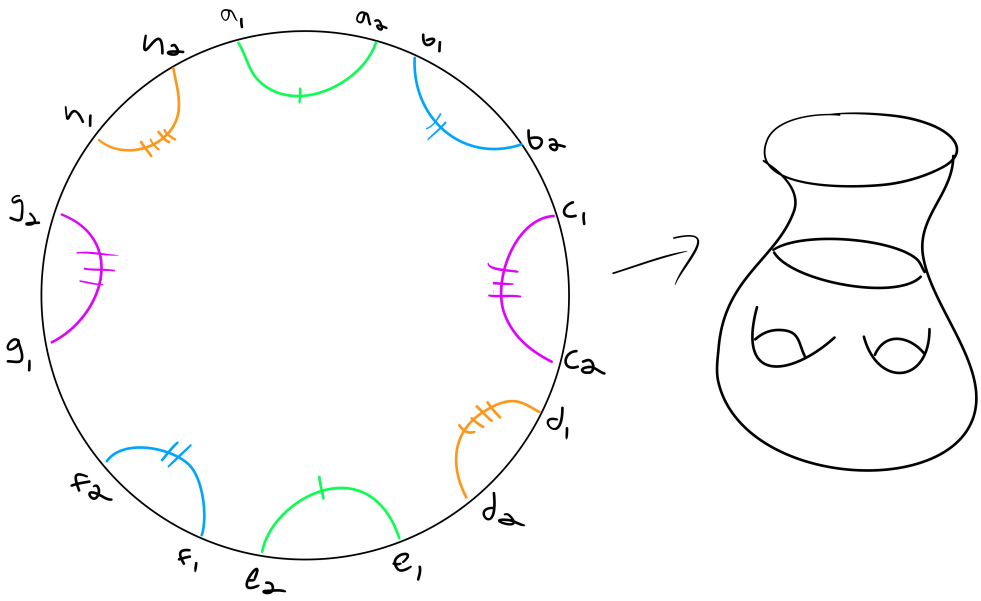

Figure 5.4: This is an example for how to glue circles to obtain a two-torus wormhole. The points a_1 and e_2 are identified, for example.

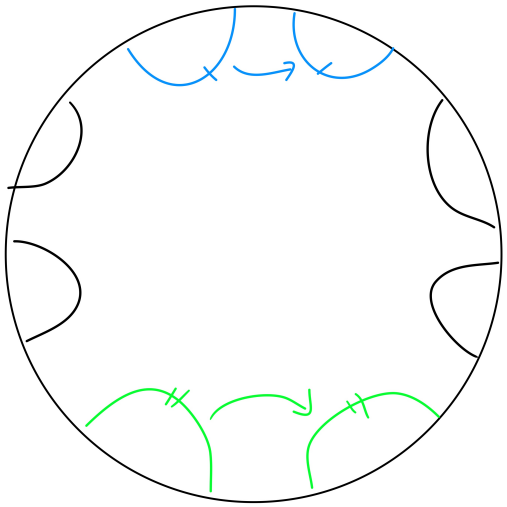

Figure 5.5: Example showing how there's a new boundary region created for every pair of neighboring circles that are identified. In this example, the blue circles and green circles are glued together, respectively. The short black boundary arc between the blue circles will become an asymptotic boundary region, as will that between the green arcs. Using the formula $G = 2h + b - 1$, we know that we are starting with $G = 4$ because there are 2×4 circles being glued together. With this choice of pairing for the blue and green circles we know that $b \geq 2$. There are only two topologically distinct choices for how to glue the black circles together. One glueing will yield the five boundary wormhole. The other glueing will yield an MBW with three boundary regions and two holes.

Chapter 6

Conclusions and Future Work

This thesis grew out of a desire to deeply understand the simplest non-trivial examples in three dimensional quantum gravity. The main contribution in this thesis was the introduction of a phase transition that emerges when the bulk has matter fields with scaling dimension less than the Hausdorff dimension of the limit set of the bulk geometry. Small progress was made towards understanding these examples but there's still work to be done. Some of the questions that we're leaving for future work include:

- To explore the bulk-boundary correspondence in quotients. The mathematical tools reviewed in Chapter 3 should be important when trying to understand how to reconstruct boundary operators given bulk data in quotients.
- To work through higher genus Ryu-Takayanagi examples, such as the two-torus wormhole from section 5.2. Especially when there's angular momentum turned on in the bulk. In this case, it would be interesting to understand the horizon dynamics.
- This work was basically a one-way flow of tools from mathematics to physics. It would be interesting to import tools going in the other direction. For example, the Ryu-Takayanagi phase diagrams in (Maxfield, 2015) oftentimes have critical points at quotients that correspond to highly symmetrical Riemann surfaces, such as the Klein quartic. The same can be said for the behavior of Rényi entropies. This yields a different way to view highly symmetrical Riemann surfaces and a way to parameterize moduli space. Another direction is that the paper (Belin, Keller, and Zadeh, 2017) hints at another way to numerically estimate Hausdorff dimension via CFT arguments. It would

also be interesting to explore whether or not entanglement inequalities have a geometric interpretation in moduli space.

- We would like to explicitly use the simultaneous uniformization theorem to analytically continue a spinning three boundary wormhole from the Lorentzian to Euclidean picture.

- Of course it would be interesting to try to extend some of these ideas to higher dimensions. There are multiple things that make this challenging, such as that there will be local degrees of freedom and that the isometry groups aren't as well studied.

- Finally, and related to the last item, we would like to study how a conformal field theory operator behaves after applying a *quasiconformal* map. Quasiconformal maps are generalizations of conformal maps, but instead of infinitesimal circles getting mapped to infinitesimal circles; circles can get mapped to elipses with a bounded amount of distortion. Quasiconformal maps play a prominent role in modern mathematics but have barely entered into the physics literature. One of the most important theorems about quasiconformal maps is the *measurable Riemann mapping theorem* which enables one to extend much of the power of the Riemann mapping theorem to higher dimensions. This is the context within which these objects may be useful in the study of conformal field theories.

Bibliography

Aminneborg, Stefan, Ingemar Bengtsson, Dieter Brill, et al. (1998). "Black holes and wormholes in (2+1)-dimensions". In: *Class. Quant. Grav.* 15, pp. 627–644. DOI: 10.1088/0264-9381/15/3/013. arXiv: gr-qc/9707036 [gr-qc].

Aminneborg, Stefan, Ingemar Bengtsson, and Soren Holst (1999). "A Spinning anti-de Sitter wormhole". In: *Class. Quant. Grav.* 16, pp. 363–382. DOI: 10.1088/0264-9381/16/2/004. arXiv: gr-qc/9805028 [gr-qc].

Balasubramanian, Vijay et al. (2014). "Multiboundary Wormholes and Holographic Entanglement". In: *Class. Quant. Grav.* 31, p. 185015. DOI: 10.1088/0264-9381/31/18/185015. arXiv: 1406.2663 [hep-th].

Banados, Maximo, Claudio Teitelboim, and Jorge Zanelli (1992). "The Black hole in three-dimensional space-time". In: *Phys. Rev. Lett.* 69, pp. 1849–1851. DOI: 10.1103/PhysRevLett.69.1849. arXiv: hep-th/9204099 [hep-th].

Barrella, Taylor et al. (2013). "Holographic entanglement beyond classical gravity". In: *JHEP* 09, p. 109. DOI: 10.1007/JHEP09(2013)109. arXiv: 1306.4682 [hep-th].

Belin, Alexandre, Ling-Yan Hung, et al. (2015). "Charged Renyi entropies and holographic superconductors". In: *JHEP* 01, p. 059. DOI: 10.1007/JHEP01(2015)059. arXiv: 1407.5630 [hep-th].

Belin, Alexandre, Christoph A. Keller, and Ida G. Zadeh (2017). "Genus Two Partition Functions and Renyi Entropies of Large c CFTs". In: arXiv: 1704.08250 [hep-th].

Belin, Alexandre and Alexander Maloney (2016). "A New Instability of the Topological black hole". In: *Class. Quant. Grav.* 33.21, p. 215003. DOI: 10.1088/0264-9381/33/21/215003. arXiv: 1412.0280 [hep-th].

Belin, Alexandre, Alexander Maloney, and Shunji Matsuura (2013). "Holographic Phases of Renyi Entropies". In: *JHEP* 12, p. 050. DOI: 10.1007/JHEP12(2013)050. arXiv: 1306.2640 [hep-th].

Bers, Lipman (1960). "Simultaneous uniformization". In: *Bulletin of the American Mathematical Society* 66.2, pp. 94–97.

Bishop, Christopher J and Peter W Jones (1997). "Hausdorff dimension and Kleinian groups". In: *Acta Mathematica* 179.1, pp. 1–39.

Bizoń, Piotr and Joanna Jałmużna (2013). "Globally regular instability of AdS_3". In: *Phys. Rev. Lett.* 111.4, p. 041102. DOI: 10.1103/PhysRevLett.111.041102. arXiv: 1306.0317 [gr-qc].

Borthwick, David (2007). *Spectral theory of infinite-area hyperbolic surfaces*. Springer.

Brill, Dieter (1998). "Black holes and wormholes in (2+1)-dimensions". In: [Lect. Notes Phys.537,143(2000)]. arXiv: gr-qc/9904083 [gr-qc].

Brill, Dieter R. (1996). "Multi - black hole geometries in (2+1)-dimensional gravity". In: *Phys. Rev.* D53, pp. 4133–4176. DOI: 10.1103/PhysRevD.53.4133. arXiv: gr-qc/9511022 [gr-qc].

Brock, Jeff (2007). "Simultaneous uniformization and limits of Kleinian groups". In:

Brown, J. David and M. Henneaux (1986). "Central Charges in the Canonical Realization of Asymptotic Symmetries: An Example from Three-Dimensional Gravity". In: *Commun. Math. Phys.* 104, pp. 207–226. DOI: 10.1007/BF01211590.

Calabrese, Pasquale and John Cardy (2009). "Entanglement entropy and conformal field theory". In: *J. Phys.* A42, p. 504005. DOI: 10.1088/1751-8113/42/50/504005. arXiv: 0905.4013 [cond-mat.stat-mech].

Calabrese, Pasquale and John L. Cardy (2004). "Entanglement entropy and quantum field theory". In: *J. Stat. Mech.* 0406, P06002. DOI: 10.1088/1742-5468/2004/06/P06002. arXiv: hep-th/0405152 [hep-th].

Calabrese, Pasquale, John Cardy, and Erik Tonni (2012). "Entanglement negativity in quantum field theory". In: *Phys. Rev. Lett.* 109, p. 130502. DOI: 10.1103/PhysRevLett.109.130502. arXiv: 1206.3092 [cond-mat.stat-mech].

Cardy, John, Alexander Maloney, and Henry Maxfield (2017). "A new handle on three-point coefficients: OPE asymptotics from genus two modular invariance". In: *JHEP* 10, p. 136. DOI: 10.1007/JHEP10(2017)136. arXiv: 1705.05855 [hep-th].

Chang, Chi-Ming and Ying-Hsuan Lin (2016). "Bootstrap, universality and horizons". In: *JHEP* 10, p. 068. DOI: 10.1007/JHEP10(2016)068. arXiv: 1604.01774 [hep-th].

Cheng, Shiu Yuen and Shing-Tung Yau (1975). "Differential equations on Riemannian manifolds and their geometric applications". In: *Communications on Pure and Applied Mathematics* 28.3, pp. 333–354.

Cho, Minjae, Scott Collier, and Xi Yin (2017a). "Genus Two Modular Bootstrap". In: arXiv: 1705.05865 [hep-th].

Cho, Minjae, Scott Collier, and Xi Yin (2017b). "Recursive Representations of Arbitrary Virasoro Conformal Blocks". In: arXiv: 1703.09805 [hep-th].

Dijkgraaf, Robbert et al. (2000). "A Black hole Farey tail". In: arXiv: hep-th/0005003 [hep-th].

Dodziuk, Jozef et al. (1987). "Estimating small eigenvalues of Riemann surfaces". In: *The legacy of Sonya Kovalevskaya (Cambridge, Mass., and Amherst, Mass., 1985), Contemp. Math* 64, pp. 93–121.

Dong, Xi et al. (2018). "Phase transitions in 3D gravity and fractal dimension". In: *Accepted for publication in JHEP; arXiv:1802.07275*.

Dorn, Harald and H. J. Otto (1994). "Two and three point functions in Liouville theory". In: *Nucl. Phys.* B429, pp. 375–388. DOI: 10.1016/0550-3213(94)00352-1. arXiv: hep-th/9403141 [hep-th].

Faulkner, Thomas (2013). "The Entanglement Renyi Entropies of Disjoint Intervals in AdS/CFT". In: arXiv: 1303.7221 [hep-th].

Gaberdiel, Matthias R., Christoph A. Keller, and Roberto Volpato (2010). "Genus Two Partition Functions of Chiral Conformal Field Theories". In: *Commun. Num. Theor. Phys.* 4, pp. 295–364. DOI: 10.4310/CNTP.2010.v4.n2.a2. arXiv: 1002.3371 [hep-th].

Giombi, Simone, Alexander Maloney, and Xi Yin (2008). "One-loop Partition Functions of 3D Gravity". In: *JHEP* 08, p. 007. DOI: 10.1088/1126-6708/2008/08/007. arXiv: 0804.1773 [hep-th].

Gubser, Steven S. (2008). "Breaking an Abelian gauge symmetry near a black hole horizon". In: *Phys. Rev.* D78, p. 065034. DOI: 10.1103/PhysRevD.78.065034. arXiv: 0801.2977 [hep-th].

Harlow, Daniel, Jonathan Maltz, and Edward Witten (2011). "Analytic Continuation of Liouville Theory". In: *JHEP* 12, p. 071. DOI: 10.1007/JHEP12(2011)071. arXiv: 1108.4417 [hep-th].

Hartman, Thomas (2013). "Entanglement Entropy at Large Central Charge". In: arXiv: 1303.6955 [hep-th].

Hartman, Thomas, Christoph A. Keller, and Bogdan Stoica (2014). "Universal Spectrum of 2d Conformal Field Theory in the Large c Limit". In: *JHEP* 09, p. 118. DOI: 10.1007/JHEP09(2014)118. arXiv: 1405.5137 [hep-th].

Hartnoll, Sean A., Christopher P. Herzog, and Gary T. Horowitz (2008). "Building a Holographic Superconductor". In: *Phys. Rev. Lett.* 101, p. 031601. DOI: 10.1103/PhysRevLett.101.031601. arXiv: 0803.3295 [hep-th].

Hawking, S. W. and Don N. Page (1983). "Thermodynamics of Black Holes in anti-De Sitter Space". In: *Commun. Math. Phys.* 87, p. 577. DOI: 10.1007/BF01208266.

Headrick, Matthew (2010). "Entanglement Renyi entropies in holographic theories". In: *Phys. Rev.* D82, p. 126010. DOI: 10.1103/PhysRevD.82.126010. arXiv: 1006.0047 [hep-th].

Headrick, Matthew et al. (2015). "Rényi entropies, the analytic bootstrap, and 3D quantum gravity at higher genus". In: *JHEP* 07, p. 059. DOI: 10.1007/JHEP07(2015)059. arXiv: 1503.07111 [hep-th].

Hou, Yong (2016). "On smooth moduli space of Riemann surfaces". In: *arXiv preprint arXiv:1610.03132*.

Keller, Christoph A., Gregoire Mathys, and Ida G. Zadeh (2017). "Bootstrapping Chiral CFTs at Genus Two". In: arXiv: 1705.05862 [hep-th].

Klebanov, Igor R. and Edward Witten (1999). "AdS / CFT correspondence and symmetry breaking". In: *Nucl. Phys.* B556, pp. 89–114. DOI: 10.1016/S0550-3213(99)00387-9. arXiv: hep-th/9905104 [hep-th].

Krasnov, Kirill (2000). "Holography and Riemann surfaces". In: *Adv. Theor. Math. Phys.* 4, pp. 929–979. arXiv: hep-th/0005106 [hep-th].

– (2002). "Analytic continuation for asymptotically AdS 3-D gravity". In: *Class. Quant. Grav.* 19, pp. 2399–2424. DOI: 10.1088/0264-9381/19/9/306. arXiv: gr-qc/0111049 [gr-qc].

– (2003). "Black hole thermodynamics and Riemann surfaces". In: *Class. Quant. Grav.* 20, pp. 2235–2250. DOI: 10.1088/0264-9381/20/11/319. arXiv: gr-qc/0302073 [gr-qc].

Kraus, Per et al. (2017). "Witten Diagrams for Torus Conformal Blocks". In: *JHEP* 09, p. 149. DOI: 10.1007/JHEP09(2017)149. arXiv: 1706.00047 [hep-th].

Lewkowycz, Aitor and Juan Maldacena (2013). "Generalized gravitational entropy". In: *JHEP* 08, p. 090. DOI: 10.1007/JHEP08(2013)090. arXiv: 1304.4926 [hep-th].

Maldacena, Juan Martin (1999). "The Large N limit of superconformal field theories and supergravity". In: *Int. J. Theor. Phys.* 38. [Adv. Theor. Math. Phys.2,231(1998)], pp. 1113–1133. DOI: 10.1023/A:1026654312961. arXiv: hep-th/9711200 [hep-th].

Maldacena, Juan Martin and Andrew Strominger (1998). "AdS(3) black holes and a stringy exclusion principle". In: *JHEP* 12, p. 005. DOI: 10.1088/1126-6708/1998/12/005. arXiv: hep-th/9804085 [hep-th].

Maloney, Alexander (2015). "Geometric Microstates for the Three Dimensional Black Hole?" In: arXiv: 1508.04079 [hep-th].

Maloney, Alexander and Edward Witten (2010). "Quantum Gravity Partition Functions in Three Dimensions". In: *JHEP* 02, p. 029. DOI: 10.1007/JHEP02(2010)029. arXiv: 0712.0155 [hep-th].

Maxfield, Henry (2015). "Entanglement entropy in three dimensional gravity". In: *JHEP* 04, p. 031. DOI: 10.1007/JHEP04(2015)031. arXiv: 1412.0687 [hep-th].

– (2017). "A view of the bulk from the worldline". In: arXiv: 1712.00885 [hep-th].

Maxfield, Henry, Simon Ross, and Benson Way (2016). "Holographic partition functions and phases for higher genus Riemann surfaces". In: *Class. Quant. Grav.* 33.12, p. 125018. DOI: 10.1088/0264-9381/33/12/125018. arXiv: 1601.00980 [hep-th].

McMullen, Curtis T (1998). "Hausdorff dimension and conformal dynamics, III: Computation of dimension". In: *American journal of mathematics*, pp. 691–721.

– (1999). "Hausdorff dimension and conformal dynamics I: Strong convergence of Kleinian groups". In:

Metlitski, Max A., Carlos A. Fuertes, and Subir Sachdev (2009). "Entanglement Entropy in the O(N) model". In: *Phys. Rev.* B80.11, p. 115122. DOI: 10.1103/PhysRevB.80.115122. arXiv: 0904.4477 [cond-mat.stat-mech].

Mumford, David, Caroline Series, and David Wright (2002). *Indra's pearls: the vision of Felix Klein*. Cambridge University Press.

Patterson, Samuel J (1976). "The limit set of a Fuchsian group". In: *Acta mathematica* 136.1, pp. 241–273.

Patterson, Samuel J, Peter A Perry, et al. (2001). "The divisor of Selberg's zeta function for Kleinian groups". In: *Duke Mathematical Journal* 106.2, pp. 321–390.

Patterson, SJ (1989). "The Selberg zeta-function of a Kleinian group". In: *Number theory, trace formulas and discrete groups*. Elsevier, pp. 409–441.

Ryu, Shinsei and Tadashi Takayanagi (2006). "Holographic derivation of entanglement entropy from AdS/CFT". In: *Phys. Rev. Lett.* 96, p. 181602. DOI: 10.1103/PhysRevLett.96.181602. arXiv: hep-th/0603001 [hep-th].

Skenderis, Kostas and Balt C. van Rees (2011). "Holography and wormholes in 2+1 dimensions". In: *Commun. Math. Phys.* 301, pp. 583–626. DOI: 10.1007/s00220-010-1163-z. arXiv: 0912.2090 [hep-th].

Strominger, Andrew (1998). "Black hole entropy from near horizon microstates". In: *JHEP* 02, p. 009. DOI: 10.1088/1126-6708/1998/02/009. arXiv: hep-th/9712251 [hep-th].

Sullivan, Dennis (1979). "The density at infinity of a discrete group of hyperbolic motions". In: *Inst. Hautes Études Sci. Publ. Math* 50.2979, pp. 171–202.

Sullivan, Dennis et al. (1987). "Related aspects of positivity in Riemannian geometry". In: *Journal of differential geometry* 25.3, pp. 327–351.

Witten, Edward (2007). "Three-Dimensional Gravity Revisited". In: arXiv: 0706.3359 [hep-th].

Yin, Xi (2008a). "On Non-handlebody Instantons in 3D Gravity". In: *JHEP* 09, p. 120. DOI: 10.1088/1126-6708/2008/09/120. arXiv: 0711.2803 [hep-th].

– (2008b). "Partition Functions of Three-Dimensional Pure Gravity". In: *Commun. Num. Theor. Phys.* 2, pp. 285–324. DOI: 10.4310/CNTP.2008.v2.n2.a1. arXiv: 0710.2129 [hep-th].

Zamolodchikov, A. B. (1984). "CONFORMAL SYMMETRY IN TWO-DIMENSIONS: AN EXPLICIT RECURRENCE FORMULA FOR THE CONFORMAL PARTIAL WAVE AMPLITUDE". In: *Commun. Math. Phys.* 96, pp. 419–422. DOI: 10.1007/BF01214585.

Zamolodchikov, Al B (1987). "Conformal symmetry in two-dimensional space: recursion representation of conformal block". In: *Theoretical and Mathematical Physics* 73.1, pp. 1088–1093.

Zamolodchikov, Alexander B. and Alexei B. Zamolodchikov (1996). "Structure constants and conformal bootstrap in Liouville field theory". In: *Nucl. Phys.* B477, pp. 577–605. DOI: 10.1016/0550-3213(96)00351-3. arXiv: hep-th/9506136 [hep-th].

Zograf, Petr Georgievich and Leon Armenovich Takhtadzhyan (1988). "On uniformization of Riemann surfaces and the Weil-Petersson metric on Teichmüller and Schottky spaces". In: *Mathematics of the USSR-Sbornik* 60.2, p. 297.