

Multiboundary Wormholes and Holographic Entanglement

Vijay Balasubramanian^{a,b}, Patrick Hayden^c, Alexander Maloney^{d,e}, Donald Marolf^f,
Simon F. Ross^g

^a*David Rittenhouse Laboratories, University of Pennsylvania
209 S 33rd Street, Philadelphia, PA 19104, USA*

^b*CUNY Graduate Center, Initiative for the Theoretical Sciences
365 Fifth Avenue, New York, NY 10016, USA*

^c*Department of Physics, Stanford University
Palo Alto, CA 94305, USA*

^d*Department of Physics, McGill University
3600 rue Université, Montreal H3A2T8, Canada*

^e*Center for the Fundamental Laws of Nature, Harvard University
Cambridge, MA 02138, USA*

^f*Department of Physics, University of California,
Santa Barbara, CA 93106, USA*

^g*Centre for Particle Theory, Department of Mathematical Sciences
Durham University, South Road, Durham DH1 3LE, UK*

Abstract

The AdS/CFT correspondence relates quantum entanglement between boundary Conformal Field Theories and geometric connections in the dual asymptotically Anti-de Sitter space-time. We consider entangled states in the n -fold tensor product of a 1+1 dimensional CFT Hilbert space defined by the Euclidean path integral over a Riemann surface with n holes. In one region of moduli space, the dual bulk state is a black hole with n asymptotically AdS₃ regions connected by a common wormhole, while in other regions the bulk fragments into disconnected components. We study the entanglement structure and compute the wave function explicitly in the puncture limit of the Riemann surface in terms of CFT n -point functions. We also use AdS minimal surfaces to measure entanglement more generally. In some regions of the moduli space the entanglement is entirely multipartite, though not of the GHZ type. However, even when the bulk is completely connected, there are regions of the moduli space in which the entanglement is instead almost entirely bipartite: significant entanglement occurs only between pairs of CFTs. We develop new tools to analyze intrinsically n -partite entanglement, and use these to show that for some wormholes with n similar sized horizons there is intrinsic entanglement between all n parties, and that the distillable entanglement between the asymptotic regions is at least $(n + 1)/2$ partite.

Contents

1	Introduction	1
2	Multiboundary black holes in 3D	4
2.1	Multiboundary solutions of 3D gravity	5
2.2	Euclidean path integral	11
2.3	Bulk phases	14
3	CFT properties of the state	17
3.1	Three-boundary state	19
3.1.1	Mapping to the sphere	19
3.1.2	Relation to OPE coefficients	22
3.1.3	Beyond the puncture limit	23
3.2	More boundaries and factorization limits	24
3.3	Reduced density matrices and factorization	27
4	Multipartite holographic entanglement	30
4.1	Holographic calculation of entanglement entropy	31
4.2	Phases of entanglement for three boundaries	33
4.3	Phases of entanglement for four boundaries	35
4.4	Entanglement versus classical correlation	37
4.5	Entanglement structure as a function of n	39
4.6	Intrinsically n -partite entanglement	40
4.7	The random state model	43
5	Conclusions	47
A	Details on Bulk Phases and the Mapping Class Group	50
A.1	Mapping Class Group	50
A.2	Spin structures	51
A.3	Hawking-Page like phase transition for three boundaries	53
B	Derivation of the k-producible bound	55

1 Introduction

Quantum systems are fundamentally distinguished from classical ones by their capacity for entanglement, a property that gives rise to many of the most counterintuitive features of quantum mechanics. A remarkable new role for entanglement has recently appeared in holographic descriptions of field theories. In the AdS/CFT correspondence, the entanglement structure of the quantum theory seems to be playing a central role in the emergence of a classical spacetime geometry.

The importance of entanglement in this context was first appreciated in the eternal black hole spacetime [1]. The eternal black hole has two asymptotic regions, connected by an Einstein-Rosen bridge. The dual state is an entangled state on the two boundaries, which can be obtained by considering the Euclidean black hole geometry. The role of entanglement in other geometries has been studied in [2, 3, 4]. Another relation between entanglement and geometry arises in the Ryu-Takayanagi (RT) formula relating the entropy of a reduced density matrix associated to a spatial subregion in a static field theory state to the area of a minimal surface in AdS anchored on the boundary of the subregion [5]. This entropy provides a measure of the amount of entanglement between the subregion and its complement. This relation has been extended to a covariant proposal [6], and the RT formula has been related to the eternal black hole construction in [7]. It has also been extended to associate a “differential entropy” with the area of certain surfaces in AdS space that do not touch the boundary [8, 9, 10, 11]. From a different perspective, it has been proposed that the general structure of the entanglement-based MERA ansatz for calculating ground state wavefunctions may provide a new understanding of the emergence of an additional dimension in the holographic description [12]. Finally, several recent efforts propose a derivation of linearized gravity from the dynamics of entanglement of the underlying quantum degrees of freedom, e.g. [13, 14, 15, 16].

All of these constructions are fundamentally bipartite. For example, the thermofield double (TFD) is a state on two CFTs dual to a spacetime connecting two asymptotic regions. Likewise, the Ryu-Takayangi formula describes entanglement between a boundary region and its complement in terms of a surface that divides the bulk space into two parts. The extensions discussed in [2, 3, 4, 8, 9, 10, 11] similarly describe entanglement between pairs of systems.

Entanglement, however, is an inherently multipartite concept. A system with many degrees of freedom can be entangled in a way that is not fully characterized by the entanglement between subsets of the degrees of freedom. An analogy can be drawn with quantum field theory, where many-point correlation functions cannot be inferred from lower-point correlations. An example of intrinsically three-party entanglement is the GHZ state of three spins: $|\text{GHZ}\rangle = (|\uparrow\uparrow\uparrow\rangle + |\downarrow\downarrow\downarrow\rangle)/\sqrt{2}$. Tracing over any one of the three spins results in a classical mixture of the product states $|\uparrow\uparrow\rangle$ and $|\downarrow\downarrow\rangle$ even though the global pure state does not factorize. Nonlocal effects can be more pronounced in the multipartite setting as well: while any state of two spins can at best violate a Bell inequality on average [17], the GHZ state can with a single measurement [18].

We wish to study multipartite entanglement in the AdS/CFT correspondence. We will focus on the AdS₃/CFT₂ case, though we expect similar considerations will apply in higher dimensions. The simplest entangled state is the TFD state, which lives in the Hilbert space $\mathcal{H} \otimes \mathcal{H}$, where \mathcal{H} is the Hilbert space of a single CFT. The TFD state is defined by the Euclidean path integral on a finite cylinder, with the two states in \mathcal{H} inserted on either end of the cylinder. The inverse temperature of the state is the conformal modulus of the cylinder. A natural multipartite generalization is the state $|\Sigma\rangle$ in $\mathcal{H}^{\otimes n}$ given by the Euclidean path integral over a Riemann surface Σ with n boundaries. The goal of this paper is to analyze the bulk connectivity and multipartite nature of entanglement in $|\Sigma\rangle$.

The state $|\Sigma\rangle$ depends on the conformal moduli of the Riemann surface Σ . For some values of the moduli, the dominant gravitational solution is a connected, multiboundary black hole: this geometry has n asymptotically AdS_3 regions connected by a wormhole with the conformal geometry of Σ . This picture was initially suggested in [1] (see also [19, 20, 21]). In this case the moduli can be interpreted as n black hole parameters (the mass, temperature or horizon area of the black hole) – one for each asymptotic region – along with a number of “internal” moduli which encode the structure of the behind-the-horizon wormhole. In other regions of moduli space the dominant bulk solution is disconnected. The transitions between these topologically distinct bulk solutions generalize the Hawking-Page phase transition.

The multipartite nature of the entanglement in $|\Sigma\rangle$ also depends on the moduli. When all boundaries are connected in the bulk through a common wormhole, one might expect multipartite entanglement to play an important role.¹ But the actual story is more complicated. In some regions of moduli space the entanglement is entirely multipartite (at leading order in the central charge), though never GHZ-like. In other regions it is entirely bipartite. Both of these behaviours are possible even in parts of moduli space where the bulk geometry is completely connected!

We begin by reviewing relevant background material in section 2. We construct multiboundary black holes as quotients of AdS_3 and describe the moduli associated with the Riemann surface Σ and the state $|\Sigma\rangle$, following [25, 26, 27, 28] and [1, 19, 20, 21]. We discuss bulk phase transitions for the states $|\Sigma\rangle$ and argue that the wormhole fragments into disconnected components in certain regions of moduli space. Black holes in the various asymptotic regions can undergo dramatic changes when a phase boundary is crossed.

In section 3 we argue that the wave function of $|\Sigma\rangle$ can be expressed as a sum of CFT n -point functions up to the action of some (complexified) conformal transformations which act separately on each individual boundary CFT. Thus $|\Sigma\rangle$ has much more structure than the simple TFD state. In the so-called puncture limit – where Σ becomes a Riemann surface with n punctures – we determine the leading conformal transformations. We then use factorization limits of the n -point functions to study disconnected phases, and identify limits of the fully-connected bulk phase where the entanglement becomes fully bipartite.

We examine the multipartite nature of the entanglement in section 4. Even in the puncture limit, the entropy of the reduced density matrices obtained by tracing over one or more CFTs is difficult to compute directly from $|\Sigma\rangle$. So instead we use the covariant HRT prescription [6], which requires only that we find the area of certain bulk extremal surfaces. Even when the bulk describes a single connected wormhole, this analysis demonstrates that the amount of entanglement and its qualitative nature both depend on the moduli: there are regimes where entanglement is largely bipartite and others where it is largely multipartite.

Unfortunately, there is no unique, agreed upon “best” measure of multipartite entanglement [29, 30, 31]. (For reviews, see [32, 33].) In the bipartite case, all pure states are asymptotically interconvertible into each other at a rate given by the ratio of their entanglement

¹One might interpret the arguments of [22] to suggest that this entanglement should resemble that of GHZ states. However, in a sense which will be made precise in section 4, we will find that GHZ-like states give universally negligible contributions to the entanglement of $|\Sigma\rangle$. This outcome might be expected from the results of [23, 24].

entropies using only “local operations and classical communication” [34]. This establishes the entanglement entropy as the essentially unique measure of bipartite entanglement. But in the multipartite case there are inequivalent forms of entanglement, some of which can be interconverted but only irreversibly. For example, any multipartite entangled state can be prepared locally and then distributed using pairwise entanglement via teleportation, but this process cannot in general be run backwards. Other forms of multipartite entanglement cannot be interconverted at all [35]. Section 4.5 and section 4.6 develop some methods to address these issues. We use these tools to show that for wormholes with n equal sized horizons the distillable entanglement between the asymptotic regions is at least $(n + 1)/2$ partite – there is none for smaller subsystems and there is always some for larger subsystems. We also show that some multiboundary black holes in which all the horizons are similarly sized must have intrinsically n -partite entanglement (when n is even) in the sense that otherwise their HRT entropies computed from minimal surface areas would be significantly smaller. When n is odd, we show that there is intrinsically $(n - 1)$ -partite entanglement.

As for the TFD state on a product of two Hilbert spaces, many entanglement properties of $|\Sigma\rangle$ match those of a simple random state model. In particular, the above results can be reproduced using only the fact that – as in [36] – if the system is divided into one large and multiple small subsystems, the small subsystems are entangled only with the large system and not with each other. Interestingly, in many cases the entropy of a given asymptotic region is determined by geometric structures behind the horizon.

We conclude with a summary and closing comments in section 5.

2 Multiboundary black holes in 3D

Three dimensional general relativity has no local gravitons, allowing one to describe rich families of solutions analytically. This is in particular the case for black holes with multiple asymptotic regions, each with a geometry asymptotic to global AdS_3 . Such spacetimes were constructed as quotients of (a subregion of) AdS_3 in [25, 26, 27]. See also [28] for the rotating case. As noted in [1] and described in detail in [21] (see also [19, 20]), they are associated with a Euclidean path integral on a certain Riemann surface Σ which provides a natural candidate for the dual CFT state $|\Sigma\rangle$. With n boundaries, the state lives in the Hilbert space $\mathcal{H}^{\otimes n}$, where \mathcal{H} is the Hilbert space of a single CFT on the cylinder.

In section 2.1 we review the bulk solutions, which are (in the non-rotating case) parameterized by a choice of Riemann surface Σ with n boundaries. The surface Σ is the spatial geometry of a constant time slice: each boundary of Σ matches on to one of the asymptotic boundaries. These solutions come in continuous families which are labelled by the moduli of the surface Σ . In section 2.2 we review the associated Euclidean path integrals which define the state $|\Sigma\rangle$. We will be somewhat brief, and refer the reader to [19, 21] for a more detailed discussion.

In section 2.3 we argue that, as the moduli are varied, there are phase transitions in which the bulk Lorentzian geometry changes topology. In particular, as one crosses a phase boundary some asymptopia disconnect from the others in the Lorentz-signature bulk geometry. In

these disconnected phases the boundaries – and thus the CFT copies – are partitioned into subsets defined by their connectivity. In these phase transitions the Lorentzian geometry changes topology, but the Euclidean geometry remains connected. The holographic construction of the wave function tells us that the mutual information between disconnected components vanishes at leading order in the central charge c . This indicates that the wave-function factorizes at this order.

Interestingly, our phase transitions changing Lorentz-signature connectivity lead to sharp changes in the entropy, energy, and temperature of *each* copy of the CFT, even if the associated boundary remains connected to many others through a Lorentzian wormhole. In particular, tuning the moduli to disconnect a single boundary from the 3-boundary wormhole leads not only to a small entropy for the disconnected CFT, but also to a sharp decrease in entropy for the two CFTs that remain connected. In the bulk, the black hole changes significantly in each asymptotic region – even in those where it does not disappear.

2.1 Multiboundary solutions of 3D gravity

AdS₃ is the Lorentzian, maximally symmetric spacetime with constant negative curvature and isometry group $SO(2, 2) \simeq SL(2, \mathbb{R}) \times SL(2, \mathbb{R})$. A global coordinate system covering all of AdS₃ is

$$\frac{ds^2}{\ell^2} = -\cosh^2 \chi d\tau^2 + d\chi^2 + \sinh^2 \chi d\phi^2. \quad (2.1)$$

The spacetime has a conformal boundary at $\chi \rightarrow \infty$. The induced metric on this boundary is conformally equivalent to the cylinder metric $d\sigma^2 = -d\tau^2 + d\phi^2$. We will be interested in solutions which are locally, but not globally, AdS₃. For Einstein gravity in three spacetime dimensions, all solutions to the equations of motion are locally AdS₃. These locally AdS₃ solutions will also be present in more complicated theories of AdS₃ gravity, though other solutions may be present as well.

A locally AdS₃ solution can be constructed by quotienting AdS₃ by a discrete subgroup Γ of its isometry group. The prototypical example is the Banados-Teitelboim-Zanelli (BTZ) black hole [37], which is the quotient of AdS₃ by the group generated by a single element γ . More complicated geometries are found by using a discrete group Γ with multiple generators [25, 26, 27, 28]. We will restrict our attention to discrete groups which lie in the diagonal $SL(2, \mathbb{R})$ subgroup of the isometry group, as only in this case is the Euclidean continuation of the geometry real. For this class of geometries, we can choose our global coordinates so that the action of Γ maps the $\tau = 0$ surface in (2.1) to itself. The full action of Γ on AdS₃ is uniquely determined by its action on this hyperbolic space.

The quotient by such a Γ is naturally described in the FRW coordinates on AdS₃:

$$\frac{ds^2}{\ell^2} = -dt^2 + \cos^2 t d\Sigma^2, \quad (2.2)$$

where $d\Sigma^2$ is the unit negative curvature metric on hyperbolic space \mathbb{H}^2 . The $t = 0$ surface in (2.2) can be taken to be the same as the $\tau = 0$ surface in global coordinates (2.1). These coordinates do not cover the whole of AdS₃, but they have the advantage that the action of

Γ on the spacetime is simply an identification on the \mathbb{H}^2 . In this coordinate patch, therefore, Γ identifies space-like separated points. The action of Γ on the full AdS_3 spacetime is more complicated, as outside this coordinate patch it might identify timelike-separated points. In order to obtain a space-time without closed timelike curves we will therefore need to remove certain portions of the full AdS_3 geometry before performing the quotient. In the present work we will not need to discuss the full spacetime geometry in detail, so we will focus on the region inside the coordinate patch (2.2).

The entirety of this patch is present without excisions in the quotient. This fact has important implications for the causal structure of our wormholes. Recall in particular that in global AdS_3 the surface $t = \pi/2$ describes the past light cone of a point, and that the generators of this cone represent null geodesics launched from the boundary at what one may call $t = 0$. As a result, given any two spatial locations on the global AdS_3 boundary, one can find timelike observers launched radially inward from those locations – at times shortly to the past of $t = 0$ – that meet in the bulk before $t = \pi/2$. This immediately implies that timelike observers who enter the quotient spacetime from distinct boundaries can meet inside our wormhole before reaching the singularity. Our wormholes thus have direct operational meaning in the bulk quantum gravity theory; they lie within the class of wormholes discussed in [38] and are distinct from the so-called long wormholes of [39, 40].

The action on AdS_3 preserves the time-reversal symmetry $t \rightarrow -t$, which will enable us to define a real Euclidean signature geometry which includes the above $t = 0$ surface as a moment of (Euclidean) time symmetry.

Technically, discrete groups Γ of this diagonal form are known as Fuchsian groups, and their action on hyperbolic space is well-studied mathematically. We restrict our attention to the case where Γ acts freely in order to avoid orbifold-type singularities. This means that Γ is generated by a collection of hyperbolic elements – elements which can, by conjugation, be put in the form $\begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}$ in $SL(2, \mathbb{R})$. Γ is then a Fuchsian group of the second kind. The quotient $\Sigma = \mathbb{H}^2/\Gamma$ is a smooth Riemann surface, possibly with boundaries. A Riemann surface of genus g with n boundaries has $6g - 6 + 3n$ moduli, and the uniformization theorem tells us any such Riemann surface can be obtained by such a quotient construction. The group Γ is isomorphic to the fundamental group $\pi_1(\Sigma)$. For each equivalence class of curves in $\pi_1(\Sigma)$ there is a unique minimum length geodesic; the length of this geodesic is just the trace $\text{Tr}(\gamma)$ of the corresponding element in Γ .

The action of Γ will have fixed points on the conformal boundary of \mathbb{H}^2 at some discrete set of points, and the action of Γ on AdS_3 has corresponding fixed points on the conformal boundary at $t = 0$. Removing from AdS_3 the causal future and past of these points yields the restricted space $\widehat{\text{AdS}}_3$ where the action of the quotient on the spacetime is free of pathologies. The spacetimes we consider are then $\widehat{\text{AdS}}_3/\Gamma$.

Two boundaries: For example, as noted above, when $\Gamma_\gamma = \mathbb{Z}$ is generated by a single element γ , the resulting spacetime is a BTZ black hole with two asymptotic regions. The black hole is static (non-rotating) since we choose to preserve time-reversal symmetry as above. The action of γ identifies a pair of geodesics in the hyperbolic plane; by choosing appropriate coordinates in the Poincaré disc representation we can take it to act as shown

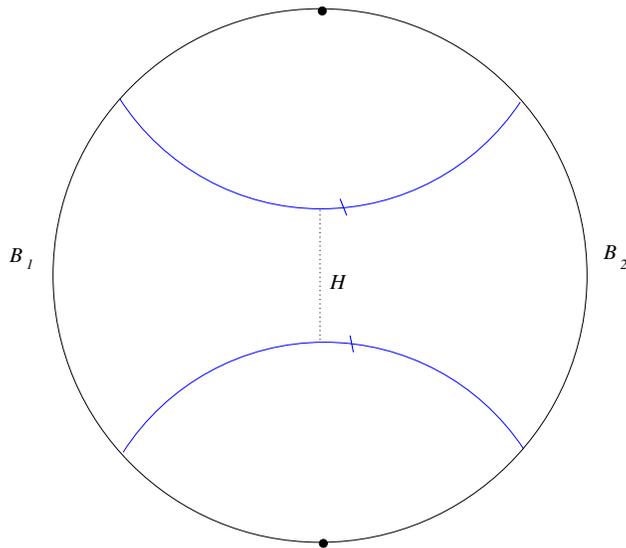


Figure 1: The $t = 0$ surface Σ in non-rotating BTZ as a quotient of the Poincaré disc. The two marked geodesics (blue in the colour version, and located symmetrically above and below the center of the figure) are identified by the action of γ . The region between them provides a fundamental domain for the quotient. B_1, B_2 become the two circular boundaries of \mathbb{H}^2/Γ . There is a minimal geodesic H , which coincides with the bifurcation surface of the BTZ event horizon; the length L of this geodesic fully characterizes the geometry of Σ .

in figure 1. The region between the two geodesics defines a fundamental domain for the quotient. The resulting surface Σ is thus topologically a cylinder $S^1 \times \mathbb{R}$ and has two conformal boundaries. This surface is the analogue of the Einstein-Rosen bridge in BTZ. As we approach either boundary the proper size of the S^1 grows, and the geometry is asymptotically \mathbb{H}^2 .

There is a unique minimal-length closed geodesic on Σ , drawn as a dotted line in figure 1. The proper length L of this minimal geodesic labels the quotient geometries uniquely.² The action of γ has two fixed points on the boundary, marked as small disks in figure 1. The restricted space $\widehat{\text{AdS}}_3$ is defined by removing from AdS_3 the future and past of these two points and the BTZ spacetime is $\widehat{\text{AdS}}_3/\Gamma$.

The minimal geodesic H is the bifurcation surface for the BTZ event horizon. In the spacetime it is the boundary of the intersection of both the causal future and the causal past of each boundary with the $t = 0$ surface, and the region to the left (right) of this geodesic provides a Cauchy surface for the domain of outer communication $I^+(\mathcal{I}) \cap I^-(\mathcal{I})$, where \mathcal{I} is the left (right) asymptotic boundary. Its length is thus related to the temperature or mass of the BTZ black hole. Indeed, for future reference we note that this black hole has temperature

$$T = \beta^{-1} = \frac{L}{4\pi^2\ell}, \quad (2.3)$$

²Genus zero is the unique exception to the general formula for the dimension of the parameter space given earlier; there is a one-parameter family of Riemann surfaces for $g = 0, n = 2$.

where L is the horizon length and (2.3) defines the inverse-temperature β . Here we have chosen the dimensionless notion of temperature associated with the dimensionless τ coordinate of (2.1) in each asymptotic region; we will also use the corresponding dimensionless notion of energy below.

The fact that the non-rotating black hole geometry is fully determined by the parameter L makes it clear that the $t = 0$ surface of any such black hole can be constructed by identifying a symmetric pair of geodesics as in figure 1. Geodesics in the Poincaré disc are circle arcs meeting the boundary at right angles. In general, in the coordinates of (2.1), a circle arc between points at $\phi = \alpha \pm \psi$ is given by

$$\tanh \chi \cos(\phi - \alpha) = \cos \psi. \quad (2.4)$$

We will describe such a circle arc as being centered at α , with opening angle ψ .

For the symmetric geodesics in figure 1 the centers are at $\alpha = \pi/2$, $\alpha' = -\pi/2$, and the opening angles are the same, $\psi = \psi'$, so we could also take ψ as the parameter characterising the identification. This is related to the length L of H by

$$L = 2\ell \tanh^{-1}(\cos \psi). \quad (2.5)$$

Note that there is a sense in which this solution contains *two* black holes, associated respectively with the right and left asymptotic regions. As we will see, this point of view will be useful in giving a unified presentation of the n -boundary cases for all n . We will therefore use this terminology below. The interesting point, however, is that for $n = 2$ one finds bulk solutions *only* when the right and left black holes are identical,³ having the same mass M and inverse temperature β .

Three boundaries: Another important case occurs when Γ is generated by two elements that each identify pairs of geodesics as depicted in figure 2. Here \mathbb{H}^2/Γ has the topology of a pair of pants. There are three circular boundaries: B_1 and B_2 are clearly circles, and the curves B_3 and B_3' connect to give a third. Note that the topology of Σ is symmetric under permutations of the boundaries; it is only our presentation that treats boundary 3 differently. The geometry of the $t = 0$ slice is asymptotically \mathbb{H}^2 at each of these boundaries, and the full spacetime will be asymptotic to global AdS₃ in each region. The spacetime in this example is thus a wormhole with three asymptotic boundaries.

The quotient has three independent minimal length geodesics, one in the homology class of each boundary, depicted as dotted lines in figure 2. The proper lengths L_a ($a = 1, 2, 3$) of these geodesics provide the three parameters labeling the geometry of Σ and hence of our spacetime.

Any pair of pants geometry can be presented with a reflection symmetry as in figure 2, with identifications between the geodesics having $\alpha_1 = \pi/2$, $\alpha'_1 = -\pi/2$, $\psi_1 = \psi'_1$, and

³This is due to our restriction to purely gravitational solutions with time-reversal symmetry and no internal topology, and to the lack of 2+1 gravitational radiation. Dynamical solutions involving matter fields can have different left and right black holes, as can solutions with a wormhole of nontrivial topology linking the two asymptotic regions.

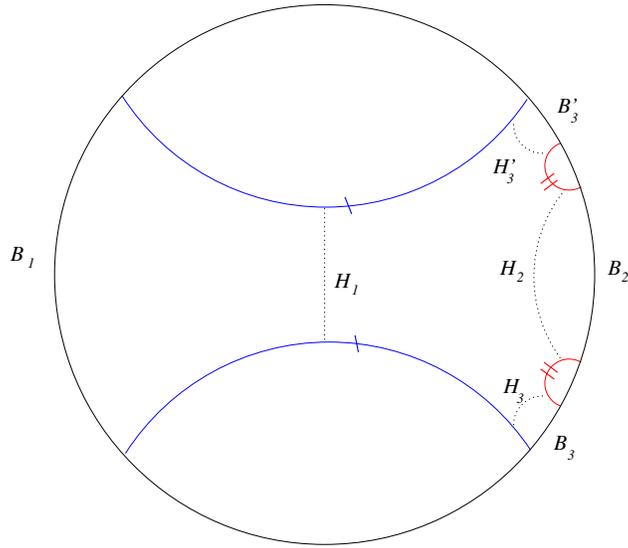


Figure 2: The $t = 0$ surface Σ in the pair of pants wormhole as a quotient of the Poincaré disc. The pairs of labeled geodesics (blue and red in colour version) are identified by the action of Γ . The region of the Poincaré disc bounded by these geodesics provides a fundamental domain for the quotient. B_1 , B_2 and $B_3 \cup B'_3$ become the desired three circular boundaries. There are corresponding minimal closed geodesics H_1 , H_2 and $H_3 \cup H'_3$, each lying at a bifurcation surface associated with past and future event horizons for the corresponding asymptotic boundary. The lengths L_a of these geodesic fully characterize the geometry of Σ .

between the geodesics having $\alpha_2 = -\alpha'_2$, $\psi_2 = \psi'_2$. We can think of α_2 , ψ_1 , ψ_2 as providing an alternative, more direct parametrization of the geometry of Σ . L_1 is related to ψ_1 as in (2.5). The calculation of the other lengths is somewhat more involved, but conceptually straightforward.

The region on our $t = 0$ surface outside any one of these minimal length geodesics is identical to a corresponding region in the $t = 0$ hypersurface of BTZ. This BTZ geometry is determined only by the horizon length L_a . This is clear in the figure for H_1 and H_2 , and we can choose a different representation of Σ to make it similarly obvious for H_3 . Thus, the region of Σ outside the geodesic again provides a Cauchy surface for the domain of outer communication attached to the corresponding asymptotic boundary, and the geometry in this region is just as in BTZ. In particular, there are exact rotation and time translation Killing fields in this region. The outgoing future and past light sheets from each minimal geodesic define future and past event horizons for observers at the corresponding asymptotic boundary.

Since the full geometry has no globally-defined Killing symmetries, the event horizons are not Killing horizons. But no observation of the geometry near the horizon can detect the difference. There is also a finite piece of the spacetime behind the horizon which is exactly BTZ. As a result, we may use the local horizon-generating Killing field near any horizon to define a notion of temperature given by (2.3). However, as we note in the next section, the natural bulk state of quantum fields on this background determined by the Euclidean continuation is not thermal even in the region outside the horizons, so this “temperature” is a property of the classical geometry only. The bulk quantum fields are expected to evolve from their non-thermal initial state towards an appropriately coarse-grained thermality [41]. But this evolution will not change the entanglement between the asymptotic regions, which are out of causal contact.

Perhaps the key difference from BTZ is the existence of a nontrivial region between the horizons H_1, H_2, H_3 . This region does not intersect the causal future or past of any asymptotic boundary. It thus lies in what is known as the “causal shadow” region of the spacetime [42]. Understanding the description of this region in terms of the boundary field theories is a particularly interesting question for holography.

This region is also interesting mathematically. We have presented a geometry on the Riemann surface with three boundaries as a subregion of the hyperbolic plane with boundaries at infinite proper distance. The hyperbolic plane is conformal to a hemisphere, so by doubling this geometry we can obtain a closed Riemann surface (the *Schottky double* of Σ) as a quotient of the sphere. This closed Riemann surface can also (and for many purposes more conveniently) be obtained as a quotient of the hyperbolic plane. (See [19, 21] for a more extensive discussion.) The representation of the Schottky double as a quotient of the hyperbolic plane is obtained by working instead with a representation of Σ where the boundaries are geodesics at finite distance. This geometry on the pair of pants is the geometry on the region bounded by the H_a in figure 2, but with some nontrivial relation between the parameters. That is, given a surface Σ specified by some geodesic lengths L_a , we can conformally map Σ to the region bounded by the H_a with some geodesic lengths $L'_a(L_a)$. Unfortunately this relation between the parametrizations is not known explicitly. We will work just with

the representation of Σ as the entire region shown in figure 2, but the representation of the pair of pants as an interior region is useful for building more complex geometries.

More boundaries: A genus g Riemann surface with n boundaries has a non-unique decomposition into $2g + n - 2$ pairs of pants, by cutting the surface along sufficient closed cycles. This decomposition provides a convenient parametrization of the moduli space of such Riemann surfaces, called the Fenchel-Nielsen parametrization, which consists of taking the 3 parameters L'_a of each of the resulting pairs of pants in the representation as a finite region in the hyperbolic plane, setting equal the lengths of boundaries that are identified, and introducing a twist parameter θ on each identified boundary. We may use this decomposition to build up the general wormhole spacetime with n boundaries by sewing together pairs of pants to build the conventional genus g surface with n boundaries and then attaching an exterior BTZ region to each of the boundaries H_a to build our wormhole.

For simplicity we will restrict our discussion to genus zero surfaces; that is, we consider just spacetimes with multiple boundaries, without introducing nontrivial topology behind the horizon.⁴ The sewn-together pairs of pants then form a tree-like structure, with no closed loops. It is reasonably straightforward to construct the full wormhole geometry as a quotient of the Poincaré disc; this is illustrated for the case of four boundaries in figure 3.

The surface Σ will have n real “external” moduli L_a , $a = 1, \dots, n$, which set the lengths of the horizons associated to each boundary. In addition, Σ will have $n - 3$ complex “internal” moduli which control the geometry of the Riemann surface once the horizon sizes are fixed. We will denote these collectively as τ_α . If we think of Σ as constructed in terms of $n - 2$ pairs of pants sewn together along some cuffs, then the τ_α are the length and twist parameters associated with each of the $n - 2$ internal cuffs. Of course, different pair-of-pants decompositions of Σ will give different coordinates τ_α on the moduli space of Σ . For example, in the $n = 4$ case, we can think of Σ as two pairs of pants, one with two external cuffs at horizons 1 and 4 and another with external cuffs at horizons 2 and 3. The internal moduli are L_{14} , the length of the internal cuff of the first pair of pants (which of course equals L_{23} , the length of the internal cuff of the second pair of pants) along with a twist parameter θ_{14} . Or one could instead choose the internal moduli to be L_{13} , the length an internal cuff of a pair of pants with external cuffs at horizons 1 and 3, along with an associated twist parameter. This just gives different coordinates on the moduli space.

2.2 Euclidean path integral

We wish to interpret the above geometries holographically, as representing bulk saddle-point descriptions of dual CFT states on their conformal boundary. In [1] it was noted that related CFT states can be constructed by a path integral over the conformal boundary of the Euclidean bulk solution. This is a simple generalization of the BTZ case (see [21] for details). Note, however, that the equivalence of the geometry near any asymptotic boundary with that of BTZ makes it natural to choose a CFT conformal frame in which the rotational

⁴Some comments on the higher genus case will appear in [43].

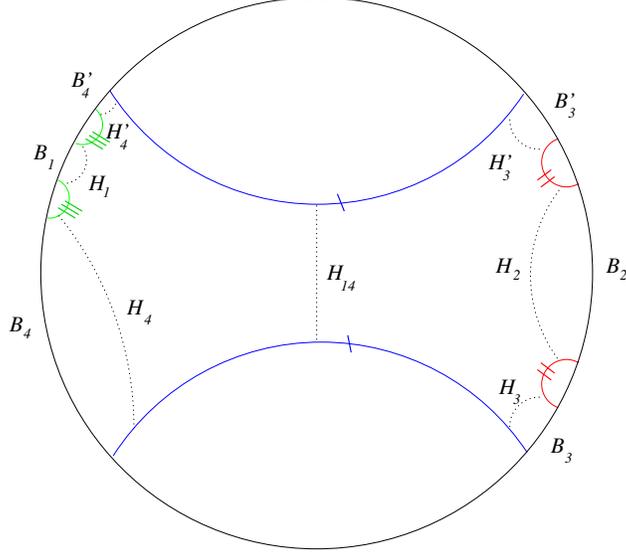


Figure 3: The $t = 0$ surface Σ in the wormhole with four boundaries as a quotient of the Poincaré disc. The pairs of marked geodesics (blue, green and red in colour version) are identified by the action of Γ . The region they bound has four asymptotic boundaries and is a fundamental domain for \mathbb{H}^2/Γ . It may be formed by sewing together two pairs of pants along H_{14} . The addition of a twist in the sewing would imply that we can no longer choose a reflection-symmetric representation of the geometry. While the first and second pair of identified geodesics are as before, the new identification introduces three new parameters; the centers α_3, α'_3 of the two identified geodesics are independent, although they can be taken to have the same opening angle, $\psi_3 = \psi'_3$. The identified geodesics are thus labeled by six parameters, corresponding to the moduli space of genus zero surfaces with four boundaries. In terms of geodesic lengths, we can take as independent parameters the lengths L_a of the four horizons, and two additional moduli characterizing the geometry of the interior region. These are naturally chosen to be the length L_{14} of the minimal geodesic H_{14} in the center and the twist θ_{14} applied along this geodesic (which is related to $\alpha_3 + \alpha'_3$). There are also similar geodesics H_{13}, H_{12} corresponding to the different ways of splitting the surface with four boundaries into two pairs of pants. But the lengths of H_{13}, H_{12} are not independent; they are determined by the moduli above.

symmetry ∂_ϕ of each CFT cylinder coincides for all time with the local rotational Killing field near the corresponding asymptotic boundary of the bulk black hole that evolves from Σ . We refer to this choice as the BTZ frame below and caution the reader that it differs from the frame chosen in [21]. Due to the conformal anomaly, this in particular leads to different expectation values for the CFT stress tensor.

On the bulk side the idea is to follow the usual Hartle-Hawking procedure and define a quantum state by a Euclidean gravitational path integral for which the wormhole spacetime describes a classical saddle-point. Since Σ represents a moment of time symmetry in the Lorentzian spacetime it has zero extrinsic curvature. Thus there is a corresponding Euclidean solution which also contains Σ as a surface of zero extrinsic curvature. The Euclidean continuation can be described as a quotient of Euclidean AdS_3 (i.e. \mathbb{H}^3) by the same discrete group Γ . The isometry group of Euclidean AdS_3 is $SO(3, 1) \simeq SL(2, \mathbb{C})$, and the Euclidean geometry will be the quotient of Euclidean AdS_3 by Γ , now regarded as a subgroup of $SL(2, \mathbb{R}) \subset SL(2, \mathbb{C})$. As in the Lorentzian case, we take the action of Γ on \mathbb{H}^3 to be determined by its action on Σ . We now see the relevance of the restriction of Γ to the diagonal subgroup; it is precisely for these groups that there is a real Euclidean geometry formed by quotienting \mathbb{H}^3 by the same discrete group.

Less abstractly, the Euclidean geometry can be constructed by analytically continuing $t \rightarrow it_E$ in (2.2) to give the Euclidean metric

$$\frac{ds^2}{\ell^2} = dt_E^2 + \cosh^2 t_E d\Sigma^2. \quad (2.6)$$

Here, as above, $d\Sigma^2$ is the constant negative curvature metric on the Riemann surface Σ . The Euclidean geometry matches onto the Lorentzian geometry at the surface of vanishing extrinsic curvature $t_E = 0$. In contrast to the Lorentzian case where (2.2) only covered a part of the geometry, this coordinate system covers the whole of the Euclidean geometry. The geometry (2.6) for all t_E is a solid handlebody whose boundary is the Schottky double of Σ . This is not entirely apparent from (2.6), which appears to have two asymptotic boundaries at $t_E \rightarrow \pm\infty$, each of which is a copy of Σ . This, however, is a quirk of the coordinate system (2.2) – it can be shown that the boundaries of the two copies of Σ are in fact glued to one another to obtain a single compact Riemann surface, the Schottky double of Σ . We refer the reader to [19, 21] for further discussion of the Euclidean multi-wormhole geometry, and [44] for a mathematical discussion of hyperbolic 3-manifolds of this type.

The suggestion of [1, 21] was to treat (2.6) for $t_E < 0$ as a saddle point of the Euclidean path integral defining our Hartle-Hawking-like state on the $t = 0$ surface of the Lorentzian geometry. If this saddle dominates the bulk path integral, then via holography it also approximates the state $|\Sigma\rangle$ defined on the boundaries B_a of Σ by the field theory path integral over the asymptotic boundary of the bulk Euclidean geometry. In the coordinates (2.6), this conformal boundary lies at $t_E \rightarrow -\infty$. It follows that its conformal geometry is identical to that of Σ . For simplicity we will describe the state $|\Sigma\rangle$ as being given by a path integral over Σ itself, though the reader should understand that we do not imply any local identification of the CFT spacetime and any particular (e.g., $t = 0$) surface in the bulk.

We conclude by noting that the full conformal boundary of (2.6) formed by joining the

two surfaces at $t_E \rightarrow \pm\infty$ generally admits no $U(1)$ isometry. As a result, there is no obvious sense in which Euclidean time is periodic, and there is no reason to expect precisely thermal behavior in any CFT copy or in any asymptotic region of the bulk.

2.3 Bulk phases

We constructed the above CFT path integral to understand the dual description of the multiboundary black hole. But we now argue that there are multiple saddle-points with the same boundary conditions, resulting in phase transitions in which different saddle points exchange dominance as the moduli of Σ are varied (see also [43]). We will keep the discussion in this section rather brief, though in appendix A.1 we give a somewhat more mathematical discussion where these phase transitions are related to the mapping class group. In a theory with fermions, there will be further interesting dependence on the choice of periodic or antiperiodic boundary conditions for the fermions on each circle. We relegate this discussion to appendix A.2 and focus here and in the main text below on the simplest case with antiperiodic boundary conditions on each boundary.

To review the two-boundary case, consider the CFT state on $\mathcal{H}_1 \otimes \mathcal{H}_2$ described by the path integral over Σ as shown in figure 1. This Σ is conformal to a cylinder, and a nice way to determine the dominant bulk saddle is to study the CFT path integral for the norm of this state. One then integrates over the torus obtained by gluing together two copies of Σ at the corresponding boundaries. We thus seek bulk saddles with torus boundaries. Such solutions to the bulk equations of motion are solid tori with a single contractible cycle. For a fixed boundary condition (fixed conformal structure of the boundary torus) there is then an infinite family of bulk saddle points, characterized by choosing a homotopy class of the boundary torus to be identified with the contractible cycle in the bulk [45].

Our torus has a \mathbb{Z}_2 symmetry exchanging the two copies of Σ that we may think of as a Euclidean time-reflection, and which implies that our state is invariant under time-reversal. Two of the above saddles preserve the time-reversal symmetry: they are called Euclidean thermal AdS (when the original CFT spatial circle is contractible) and the Euclidean BTZ black hole (when the Euclidean time direction is contractible). Here the term thermal is associated with having a Euclidean time translation symmetry, so that the period β of this Euclidean time defines an inverse temperature. Taking the spatial circle to have period 2π , there is a Hawking-Page phase transition $\beta = 2\pi$, with Euclidean thermal AdS₃ dominating for small temperatures (large β) and Euclidean BTZ dominating for high temperatures. At least one of these always dominates over those bulk geometries that break time-reversal symmetry.

Since Euclidean time-reversal symmetry implies that the $t = 0$ surface has zero extrinsic curvature, we can use the bulk geometry on this surface as initial data to build a unique Lorentz-signature spacetime. This spacetime is just the Wick rotation of the Euclidean saddle. In this sense it is convenient to speak of our bulk path integral being dominated by the corresponding Lorentz-signature spacetime which in our case is either two copies of global AdS₃ or BTZ as implied by the terminology above. In the global AdS₃ case the quantum fields on the two copies combine to form a thermofield double state.

From the Lorentzian point of view, the thermal AdS and Euclidean BTZ saddles give contributions to the wave function which described different Lorentzian space-times. The state is obtained by matching onto a Lorentzian signature geometry at a moment of time reflection symmetry (i.e. a surface of vanishing extrinsic curvature). In the thermal AdS saddle, the surface of vanishing extrinsic curvature is a pair of disconnected discs; this geometry constructs the thermofield double entangled state between a pair of disconnected AdS space times. In the Euclidean BTZ saddle the surface of vanishing extrinsic curvature is a cylinder with two asymptotic regions; this geometry constructs the BTZ black hole with two boundaries connected by a wormhole.

The three-boundary case has a similar but richer structure. In the case where Σ is a pair of pants, we can glue together two copies of Σ to obtain a genus two Riemann surface. The bulk solution of interest is a handlebody, the three-manifold constructed by filling in this Riemann surface. The Riemann surface has four elementary cycles, of which two are contractible in the bulk. Different Euclidean saddle points can then be constructed by specifying which combinations of cycles are contractible in the bulk. We are interested in those choices which have a \mathbb{Z}_2 time-reversal symmetry, so that they can be analytically continued to real Lorentzian geometries. There are three different choices of geometry, shown in figure 4, which have this symmetry. These choices correspond to three different ways one can construct a genus two Riemann surface by gluing together two pairs-of-pants.

As in the genus one case, each Euclidean geometry is regarded as preparing a state in the Lorentzian theory. To understand the Lorentzian interpretation, we need to find the surface of vanishing extrinsic curvature in each case. In the first case the surface of vanishing extrinsic curvature is three disconnected discs. Thus this geometry constructs an entangled state in three different disconnected AdS space-times. In the second case the surface is the pair of pants Σ . So this geometry constructs a three-boundary (pair of pants) wormhole connecting all the boundaries. In the third case the surface of zero extrinsic curvature is topologically a disc and a cylinder. This geometry constructs a BTZ black hole connecting one pair of boundaries which is entangled with one copy of global AdS₃. In fact, this case yields three distinct saddles depending on which pair of boundaries are taken to be connected.

We would expect varying moduli (parametrized by the horizon lengths L_a) to induce transitions between these different phases. In the limit where all L_a are large one expects the three-boundary wormhole to dominate; for all L_a small one expects the 3-global-AdS₃ phase to dominate; while for one L_a small with the other two large we expect the global-AdS₃-plus-BTZ phase to dominate.

By construction our $\beta_a = \frac{4\pi^2\ell}{L_a}$ are the inverse temperatures of the black holes in the connected 3-boundary wormhole phase. But they will not generally be the inverse temperatures in the BTZ-plus-global AdS₃ phase. This is because the $t = 0$ geometry outside a given bulk BTZ horizon after disconnection is not generally the same as that in the associated asymptotic region of the fully-connected bulk wormhole of Σ . This is clear from the fact that BTZ solutions require identical black holes on right and left, while the fully-connected wormhole allows all β_a to be distinct. Our phase transition will thus generically be associated with discontinuous changes in the geometry of the exterior black hole regions, even when the black hole does not completely disappear. In fact, we will argue in section 3.3 below that

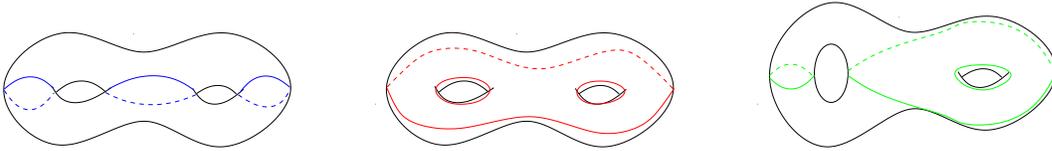


Figure 4: Different decompositions of the genus two Riemann surface which preserve a \mathbb{Z}_2 reflection symmetry. In each case the boundary of the handlebody is decomposed into two pairs of pants. If we take the pair of pants Σ to have the same geometry, the geometry of the resulting genus two surface is different in each case, but the different decompositions describe the same family of bulk handlebody geometries up to diffeomorphism. In the bulk handlebody, the cycles indicated in the left figure are contractible, while those indicated in the second are not. Slicing the bulk handlebody along the cycles in the left figure thus provides a bulk initial data surface consisting of three disconnected discs, corresponding to three copies of global AdS. The second slicing produces a bulk initial data surface for a three-boundary wormhole. The final slicing, along the cycles indicated in the right figure, yields a bulk initial data surface consisting of an annulus and a disconnected disc, corresponding to BTZ and a disconnected copy of global AdS.

the black hole temperature changes by a factor of two even when e.g. $\beta_1 = \beta_2$ and the phase transition occurs due to varying β_3 .

Further, while inverse temperatures are a useful way to characterize the geometries of the saddle-points that include black holes, we will see in Sec. 3 that the quantum state of matter fields is generally non-thermal even when restricted to the exterior regions. In particular, while the bulk quantum fields on the global AdS₃ retain some order-one entanglement with the other spaces, the associated density matrix on one such factor need not be precisely thermal. In the next section we will often focus on the puncture limit of small L_a ; we will find that in this limit at least there is some thermal behaviour in the CFT state, with the contribution of high-energy states having a Boltzmann suppression determined by the β_a .

In addition to the above saddle-points obtained by various slicings of the handlebody, there will also be non-handlebody solutions in the bulk, as noted in [46]. These could in principle produce additional phases. However, it was conjectured in [46] that such non-handlebody saddles are always subleading in the path integral compared to the handlebodies. We will assume that this is true, and not consider the non-handlebody geometries further.

The phase transitions can in principle be calculated by evaluating the bulk action for the different saddles. But in special subspaces of the moduli space, we can avoid doing this explicit calculation. In the two-boundary case, the transition between two copies of thermal AdS₃ and BTZ occurs when the boundary torus is square, that is the length of the time and space circles are equal. This is a point of enhanced symmetry in the torus moduli space. Phase transitions in the moduli space of genus two surfaces should also occur at points of enhanced symmetry. For example, the subspace where all the temperatures are equal includes the Bolza surface, the most symmetric genus two surface. There should then be a transition from the three AdS₃ saddle to the three-boundary wormhole saddle when we reach the Bolza surface. In appendix A.3 we show that this occurs at β of order one,

as expected. By continuity, we can then argue that when the β_a are of similar size and all small, the dominant saddle should be three copies of AdS_3 , while when they are of similar size and all large, it should be the three-boundary wormhole.

For $n > 3$ boundaries there will be a similar structure but with more possibilities. Relevant phases include the n -boundary connected wormhole as well as any combination of disconnected components i having respectively $m_i \geq 1$ boundaries with $\sum_i m_i = n$.

Understanding the details of this phase structure is interesting in its own right, but the important point for our present discussion is that the CFT states we consider do not always correspond to a single connected wormhole in the bulk. In particular, much of the discussion in the next section will be carried out in a puncture limit where the lengths L_a are small so that – at least for the simplest choice of spin structure – fully connected bulk saddles will not dominate. However, as in the two-boundary case, some of the qualitative features of the entanglement in the CFT state will be the same in any phase.

3 CFT properties of the state

When there are n asymptotic regions in the bulk Lorentzian spacetime, the dual CFT state $|\Sigma\rangle$ is an element of $\mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots \mathcal{H}_n$ where \mathcal{H}_i is the Hilbert space of the CFT on a circle. We are interested in a state $|\Sigma\rangle$ defined by the CFT path integral over a Riemann surface Σ with n boundaries. Holographically, this Σ also gives the asymptotic boundary conditions of the corresponding bulk path integral, and in the particular case when the bulk is dominated by the fully-connected wormhole we saw in section 2.2 that Σ is also the geometry of the bulk $t = 0$ surface.

In the CFT path integral, if states $|\phi_1\rangle \in \mathcal{H}_1, \dots, |\phi_n\rangle \in \mathcal{H}_n$ are used to fix boundary conditions on the respective boundaries of Σ , this path integral computes the amplitude $\langle \Sigma | \phi_1 \phi_2 \dots \phi_n \rangle$. Working in a basis of energy eigenfunctions $|i\rangle$ (of energy E_i) in each \mathcal{H}_a we may write the state as

$$|\Sigma\rangle = \sum_{i_1 \dots i_n} A_{i_1 \dots i_n} |i_1\rangle_1 \cdots |i_n\rangle_n, \quad (3.1)$$

where the coefficients $A_{i_1 \dots i_n}$ are functions of the $3n - 6$ moduli of Σ , determined by the Euclidean path integral on Σ . For simplicity we again restrict attention to the case where Σ has genus zero.

In the two-boundary case the $A_{i_1 i_2}$ can be determined explicitly. There Σ is conformal to a flat cylinder, and the BTZ choice of conformal frame makes (3.1) invariant under both time-reversal and the natural rotational symmetry. If its circumference is 2π , the cylinder has length $\beta/2$, where β is the periodicity of the dimensionless bulk imaginary time coordinate corresponding asymptotically to the Wick rotation of τ in (2.1). Thus the CFT path integral produces the usual thermofield-double state

$$|\Sigma\rangle = \sum_i e^{-\beta E_i/2\ell} |i\rangle_1 |i\rangle_2 \quad (3.2)$$

where β coincides with the inverse temperature one would assign to both the BTZ and double-global- AdS_3 bulk saddles.

In the multiboundary case the path integral over Σ is more complicated and we will not generally be able to write the coefficients explicitly. In general the $A_{i_1\dots i_n}$ will be non-vanishing even if the indices i_a are unequal, unlike the two-boundary case. Our state $|\Sigma\rangle$ is thus more complicated than a GHZ-like state, in that in the GHZ state $|\text{GHZ}\rangle = (|\uparrow\uparrow\uparrow\rangle + |\downarrow\downarrow\downarrow\rangle)/\sqrt{2}$ only the diagonal elements $A_{\uparrow\uparrow\uparrow}$ and $A_{\downarrow\downarrow\downarrow}$ are non-vanishing.

To evaluate (3.1) we will pursue the following general strategy. First the Riemann surface Σ can always be conformally mapped to a sphere with n holes h_a . Using the state-operator correspondence, we can map the states $|i\rangle_a$ fixing the boundary conditions on these holes to local insertions of the corresponding operators O_a at some point inside each hole; the $A_{i_1\dots i_n}$ are thus related to the n -point correlation functions $\langle O_1 \dots O_n \rangle$. The mapping from states to operators will introduce a factor depending on the size and shape of the hole h_a . As a result, the $A_{i_1\dots i_n}$ are related to the $\langle O_1 \dots O_n \rangle$ by the action of operators V_a acting on \mathcal{H}_a . These V_a are essentially complexified conformal transformations that depend only on the moduli of the Riemann surface (L_a, τ_a in the notation of the previous section) and so are universal for all CFTs. In this way, the coefficients $A_{i_1\dots i_n}$ can be written as a product of a universal piece that depends on the moduli of the Riemann surface (and hence on the parameters of the Lorentzian spacetime), and a piece that depends on the specific field content of the CFT via the n -point functions on the sphere.

We will be able to carry out this procedure explicitly when the Riemann surface Σ has genus zero and is close to the puncture limit ($L_a \rightarrow 0$). In this limit we will see that we can choose a conformal frame where the shrinking holes are round; the moduli of Σ then correspond to radii and locations of the holes. Using this picture we will show that, up to exponentially suppressed corrections and a convention-dependent rotation,

$$V_a = \exp\left(-\frac{1}{2}\hat{\beta}_a H_a\right), \quad \text{with } \hat{\beta}_a = \beta_a - \beta_a^0. \quad (3.3)$$

Here H_a is the Hamiltonian on boundary a , and $\beta_a = \frac{4\pi^2\ell}{L_a}$ is the inverse temperature of the BTZ geometry in the exterior region of the surface Σ (i.e. the region between the minimal geodesic whose length is L_a and the boundary a). The parameter β_a^0 depends only on the moduli that are held finite in this limit.⁵ Thus, at least in the puncture limit, the $A_{i_1\dots i_n}$ involve a Boltzmann suppression similar in character to the two-boundary case.

A key simplification in the puncture limit is that the V_a (3.3) are diagonal in the energy basis, so that the coefficients $A_{i_1\dots i_n}$ become simply proportional to the n -point functions. This simple relation holds only in the puncture limit; more generally we do not expect the energy basis to diagonalize the V_a .

In section 3.1 we carry out this puncture limit calculation in detail for the three-boundary case, and comment on the extension beyond the puncture limit. In section 3.2 and 3.3 we study interesting limits in the general n -boundary moduli space. In particular, limits in which n -point functions factorize lead to corresponding factorization of the state (3.1). A similar discussion at the level of density matrices will then allow us to make further observations regarding the entanglement structure of (3.1).

⁵There is also a more general puncture limit where only some L_a are taken small. The corresponding V_a again take the form (3.3), but the β_a^0 can then also depend on the L_b that remain finite.

3.1 Three-boundary state

We now explore the relation between the CFT three-point function and the state on three boundaries. Because there are no internal moduli, this is the simplest example of the connection to n -point functions. For three boundaries, the CFT state $|\Sigma\rangle$ is an element of $\mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \mathcal{H}_3$, where \mathcal{H}_a is the Hilbert space of the CFT on a circle. Introducing a basis of energy eigenstates $|i\rangle_a$ on each boundary, it can in general be written

$$|\Sigma\rangle = \sum_{ijk} A_{ijk} |i\rangle_1 |j\rangle_2 |k\rangle_3, \quad (3.4)$$

where A_{ijk} is a set of coefficients whose structure we wish to determine. These coefficients are defined by a path integral over some pair of pants geometry, with the specified states on each boundary.

In section 2 we parametrized the three-dimensional moduli space of pair-of-pants metrics by the lengths L_1, L_2, L_3 of the minimal geodesics. The A_{ijk} will be functions of these parameters. In a theory with fermions, there will also be significant dependence on the choice of periodic or antiperiodic boundary conditions for the fermions on each circle; we relegate this discussion to appendix A.2.

3.1.1 Mapping to the sphere

As outlined above, we begin by conformally mapping our problem to a path integral over the round sphere with three holes h_a . We will do explicit calculations for small L_a . At least in some open set in which all L_a are small, we will be able to choose a conformal frame in which the holes are round holes of circumference $2\pi r_a$ centered at equally spaced points around the equator; the moduli of Σ are mapped to the three radii r_a . Far from the puncture limit, such a description may not be possible because the conformal transformations mapping the individual holes to round circles can produce overlapping circles, though it should be possible to understand the description for general L_a by a suitable analytic continuation.

We do not know the explicit conformal map from Σ to the sphere even for small L_a . But in the regime of small L_a , we can determine the radii r_a as a function of the L_a up to a universal factor. We do so by breaking the problem up into two overlapping regions: the exterior legs of Σ , which have the geometry of the $t = 0$ surface in BTZ, and an interior region I , which is well-approximated by a universal geometry for small L_a .

In the limit of vanishing L_a , the geometry of Σ approaches a universal geometry Σ_0 . This limit is illustrated in figure 5, which redraws figure 2 to make it clear that in this limit all minimal geodesics recede to infinite proper distance from the interior region I , and the boundary segments B_a shrink to points.⁶ For small L_a , the geometry in an interior region I will be close to that of Σ_0 , and we may approximate the map from I to the sphere by that from the corresponding region of Σ_0 to the sphere. (This approximation breaks down before we reach the minimal geodesics H_a .)

⁶In the limit, B_3 and B'_3 approach points $\pm\phi_0$ on the boundary, which may be different for different values of L_a , but different values of ϕ_0 are related by an $SL(2, \mathbb{R})$ transformation of the disc, so they define equivalent geometries, and there is a unique limiting geometry Σ_0 .

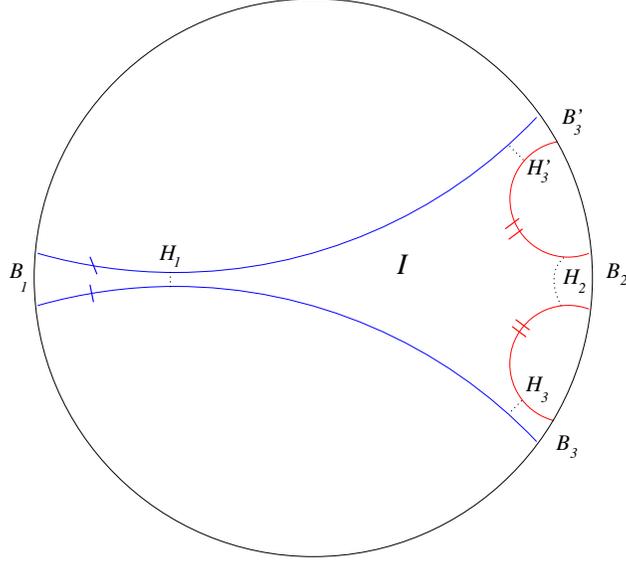


Figure 5: The limit of small L_a for the surface Σ , showing the region I between the minimal geodesics. Here the identifications are drawn in a different presentation to make it obvious that all the minimal geodesics go off to infinite proper distance in the limit of small L_a .

The universal geometry Σ_0 is a hyperbolic metric with three infinitely long cuspy arms. There is some conformal map from this surface to a punctured sphere, sending the end of each cusp to one of the punctures, which we can take to be equally spaced on the equator of the sphere. We do not know the full map from Σ_0 to the sphere, but it is easy to give its form near one of these punctures. As noted above, for any Σ the geometry around each of the minimal geodesics is locally that of the $t = 0$ surface for a BTZ black hole. The $L_a \rightarrow 0$ limit gives the zero mass black hole which is described by the hyperbolic cusp geometry

$$ds_{cusp}^2 = \ell^2 \frac{d\rho^2}{\rho^2} + \rho^2 d\phi^2. \quad (3.5)$$

Thus, the geometry near the cusp at the end of each arm of Σ_0 (see Fig. 5) will be isometric to (3.5) in an open set containing the end of each arm. In (3.5) the coordinate ϕ has period 2π and $\rho \in (0, \infty)$ with $\rho \rightarrow 0$ at the tip of the cusp.

We may take the conformal map from Σ_0 to our punctured sphere to have the property that $\rho = \text{constant}$ curves near the tip of each cusp map to closed curves on the sphere that become round circles as $\rho \rightarrow 0$. This is clearly the case when we conformally map the precise cusp geometry (3.5) to a disk (even when the disk is considered as a part of a sphere; i.e., with Ricci scalar $R_{disk} = 2$). The conformal mapping from Σ_0 (which interpolates between three cusps) to a sphere is more complicated, but is fundamentally governed by the elliptic differential equation obtained from the Weyl transformation properties of the Ricci scalar. In particular, for $ds_{sphere}^2 = \Omega^2 ds_{\Sigma}^2$ we require

$$\Omega^2 R_{sphere} = R_{\Sigma} + 2\nabla^2 \ln \Omega. \quad (3.6)$$

Since this equation approaches that governing the above cusp-to-disk conformal map near the tip of each cusp, we may seek a solution of (3.6) that approximates the cusp-to-disk map near the tip of each cusp. Thus, near one of the punctures, the map from the region near the cusp to the punctured sphere will be

$$r \approx r_d e^{-\ell/\rho}, \quad (3.7)$$

where r is the radial distance from the puncture on the sphere and r_d represents the zero mode of Ω (i.e. the overall scale of the map). This degree of freedom is not fixed by the above discussion in the cusp region and encodes our ignorance of the full map from Σ_0 to the sphere.

Corrections to the cusp-to-disk map implied by (3.6) may further be computed perturbatively in the cusp region. Perhaps the easiest way to understand the nature of these corrections is to work in the sphere conformal frame and to recall that the Green's function for Laplace's equation on a disk is of the form r^m near the origin where m is the angular momentum of the desired source and r is the usual radial coordinate on the disk. Effects that break rotational symmetry thus decay at least as r as $r \rightarrow 0$. These corrections are thus exponentially suppressed for $\rho \ll \ell$.

For small L_a , the surface Σ has an interior region I which is well-approximated by a corresponding region of Σ_0 , which extends down into the cusp regions where the approximate map (3.7) is valid. The conformal map to the sphere for this region should be well-approximated by the map from Σ_0 to the sphere. The remaining problem is then to map the regions of Σ outside I into these holes. This can be done using the fact that the geometry on each exterior arm of Σ is just that of BTZ at $t = 0$. Recall that this description is valid even a finite distance inside the horizon. In fact, as should be clear from figure 5, in the limit of small L_a it remains valid *very* far inside the horizon. The metric in such a region is⁷

$$ds_{btz}^2 = \ell^2 \frac{d\rho^2}{\rho^2 + \rho_a^2} + (\rho^2 + \rho_a^2) d\phi^2 \quad (3.8)$$

with $\rho_a = L_a/2\pi$. Here we again take ϕ to have period 2π but now $\rho \in (-\infty, \infty)$. Since L_a is small, the merger of this arm with the others takes place essentially as in Σ_0 . The lack of free parameters in Σ_0 means that, in the coordinates of (3.5),(3.8) this merger must take place at some ρ of order ℓ . Thus (3.8) is the exact metric on Σ in the region $-\infty < \rho \lesssim \ell$, where $\rho = -\infty$ corresponds to the exterior boundary and $\rho = 0$ is the minimal geodesic.

The region covered by the BTZ coordinates (3.8) overlaps with I , so we can determine the conformal map to the sphere in the region $\rho \ll \ell$ by requiring agreement with (3.7) in the region $\rho_a \ll \rho \ll \ell$, where the metric (3.8) is well approximated by (3.5). The map from Σ to the sphere in this region should coincide to leading order in ρ_a/ρ with that from Σ_0 to the sphere, and map the rotational Killing field near the a th boundary of Σ to that of the sphere about the associated center c_a to leading order in ρ/ℓ . Using a radial coordinate around the center c_a it must then take the form

$$r = r_d e^{(\ell/\rho_a)[\tan^{-1}(\rho/\rho_a) - \pi/2]} \quad (3.9)$$

⁷The metric (3.8) is equivalent to that of [37] when written in terms of $r_{BTZ} = \sqrt{\rho^2 + \rho_a^2}$, but we have chosen the form (3.8) to make manifest the agreement with (3.5) for $\rho \gg L_a$.

to leading order in $\frac{\ell}{L_a}$, $\frac{\rho_a}{\ell}$. The asymptotic boundary $\rho = -\infty$ thus maps to

$$r_a = r_d e^{-\pi\ell/\rho_a} = r_d e^{-\frac{1}{2}\beta_a}, \quad (3.10)$$

where $\beta_a = \frac{4\pi^2\ell}{L_a}$ is the inverse temperature associated to the BTZ geometry (3.8). Recall that r_d represents an overall scale in the map from the limiting geometry Σ_0 to the sphere, and is hence independent of the moduli of Σ . In cases with more boundaries the analogous factors can depend on the internal moduli, though they are again independent of any horizon lengths which become small.

Thus in the puncture limit we can conformally map each hole to a round hole of radius r_a . Taking the state-operator map to be that given by the path integral over the unit-radius disk in the plane, this implies that the operator V_a is defined by the path integral over an annulus between the hole of radius r_a and the unit circle. This yields $V_a = e^{-\frac{1}{2}\hat{\beta}_a H_a}$ as claimed in (3.3), with

$$\hat{\beta}_a = \beta_a - 2 \ln r_d. \quad (3.11)$$

Each V_a in principle also includes a convention-dependent rotation associated with how one fixes the location of $\phi = 0$ in each asymptotic region. However, this rotation can be set to zero by noting that Σ in figure 2 has a \mathbb{Z}_2 isometry that acts as reflection on each asymptotic boundary. In the sphere conformal frame and with the punctures chosen to lie on the equator, this symmetry exchanges the northern and southern hemispheres. In either frame the symmetry may be used to define a coordinate ϕ near each boundary/puncture such that it locally acts as $\phi \rightarrow -\phi$. In particular, the state-operator map at c_1 is related to that used at c_2, c_3 by a rotation of the equator about the poles. This fixes compatible conventions on Σ and the sphere that sets this extra rotation to zero.

3.1.2 Relation to OPE coefficients

We have shown above that to leading order in small L_a our state takes the form

$$|\Sigma\rangle = \sum_{ijk} \langle O_i(c_1) O_j(c_2) O_k(c_3) \rangle V_1 V_2 V_3 |i\rangle_1 |j\rangle_2 |k\rangle_3, \quad (3.12)$$

where $\langle O_i(c_1) O_j(c_2) O_k(c_3) \rangle$ denotes the three-point function on the sphere with the c_a equally spaced around the equator and the V_a are given by (3.3) and (3.11). Many readers will find it useful to rewrite (3.12) in terms of the standard three-point function on the plane. Since the operators are not generally scalars, the full expression on the plane takes the form

$$\langle O_i(z_1) O_j(z_2) O_k(z_3) \rangle_{\mathbb{R}^2} = \frac{C_{ijk}}{(z_1 - z_2)^{\Delta_{ij}/2} (z_1 - z_3)^{\Delta_{ik}/2} (z_2 - z_3)^{\Delta_{jk}/2} (\bar{z}_1 - \bar{z}_2)^{\tilde{\Delta}_{ij}/2} (\bar{z}_1 - \bar{z}_3)^{\tilde{\Delta}_{ik}/2} (\bar{z}_2 - \bar{z}_3)^{\tilde{\Delta}_{jk}/2}}, \quad (3.13)$$

where $\Delta_{ij} = \Delta_i + \Delta_j - \Delta_k$ etc, the $\Delta_i, \tilde{\Delta}_i$ are the left- and right-moving conformal weights of the operators O_i , corresponding to $E_i \pm J_i$ in terms of the energies E and the angular

momenta J of the states $|i\rangle$ in the CFT on the cylinder of circumference 2π . The C_{ijk} are the OPE coefficients, and the z_a are the images on the complex plane of the c_a on the sphere. On the plane, it is natural to take $z_1 = 1$, $z_2 = e^{2\pi i/3}$, $z_3 = e^{-2\pi i/3}$ and to define the state-operator map at each of z_1, z_2, z_3 to be related by translations of the plane⁸. This is reflected in the fact that, using the rotation invariance of the vacuum state to set $J_i + J_j + J_k = 0$ for any nonzero three-point function, (3.14) at our z_1, z_2, z_3 becomes

$$\langle O_i(z_1)O_j(z_2)O_k(z_3) \rangle_{\mathbb{R}^2} = C_{ijk} e^{i\frac{2\pi}{3}J_j} e^{-i\frac{2\pi}{3}J_k} e^{-\ln 3(E_i+E_j+E_k)}. \quad (3.14)$$

Here we have assumed that J_a has integer eigenvalues corresponding to our focus on antiperiodic boundary conditions for fermions. The more general situation involves branch cuts.

Using the same conventions on Σ as above, with the coordinate ϕ being fixed by taking the reflection to be $\phi \rightarrow -\phi$, the factor of $e^{i\frac{2\pi}{3}J_j - i\frac{2\pi}{3}J_k}$ in (3.14) precisely cancels the rotations in the V_a . Thus, we can write the three-boundary state as

$$|\Sigma\rangle = \sum_{ijk} C_{ijk} e^{-\frac{1}{2}\tilde{\beta}_1 H_1} e^{-\frac{1}{2}\tilde{\beta}_2 H_2} e^{-\frac{1}{2}\tilde{\beta}_3 H_3} |i\rangle_1 |j\rangle_2 |k\rangle_3, \quad (3.15)$$

where

$$\tilde{\beta}_a = \beta_a - 2 \ln r_d - 2 \ln 3. \quad (3.16)$$

The three-boundary state in the puncture limit is thus determined by the CFT data – the OPE coefficients C_{ijk} – together with universal factors which carry all moduli dependence. The latter give a Boltzmann suppression of a similar character to that in the thermofield double state.

3.1.3 Beyond the puncture limit

It is clear that similar asymptotic forms are obtained in other limits where some subset of the L_a 's vanish with the others held fixed. We will also see later in section 3.2 that analogous behavior occurs in the limit $L_3 \rightarrow \infty$ with $L_1, L_2 \rightarrow 0$. But we may also relate the three-boundary state to the three-point function for more general L_a .

Even when the holes h_a are not round, we can write the state as in (3.12) with operators V_a defined by a path integral over an annulus cut out of the complex plane whenever the h_a are small enough that the required disks do not overlap on the sphere⁹. The outer boundary is a circle of unit radius, and the inner boundary is defined by conformally mapping h_a to the plane. We may now use the result that any annulus is conformal to an annulus with two round boundaries (see e.g. [47]) to write the operator V_a in terms of some $\exp -\frac{1}{2}\hat{\beta}H_a$. However, we must also take into account how this conformal transformation acts on the states that we attach to this operator. States are of course functionals of field configurations on the

⁸In contrast, the convention we chose on the sphere would induce state-operator maps at z_1, z_2, z_3 on the plane that are related by rotations about the origin.

⁹We expect that this construction may be analytically continued as functions of the L_a to cases where these disks would in fact overlap. In this sense it can be used for general L_a .

boundary; i.e., functionals of field configurations $\Phi(\phi)$ written in terms of some coordinate ϕ on the S^1 . Changing the choice of this ϕ enacts a diffeomorphism of the circle, which is the diagonal subgroup of the conformal group on the cylinder that acts simultaneously on both left- and right-movers. Since mapping our original annulus to, say, a round cylinder of circumference 2π will generally involve a Weyl factor Ω that varies along both boundaries, and since states are naturally attached to both the cylinder path integral and the path integral on Σ using coordinates ϕ associated with proper distance along each boundary, we see that we must generally allow for the action of such a (real) S^1 -diffeomorphism at each boundary. In other words, for general L_a any operator V_a will take the form

$$V_a = U(h_a)e^{-\frac{1}{2}\hat{\beta}H_a}\tilde{U}(h_a), \quad (3.17)$$

where U, \tilde{U} may be written

$$U = \exp\left(i\sum_{n\geq 0} a_{I,n}(L_n + L_{-n} - \bar{L}_n - \bar{L}_{-n})\right), \quad \tilde{U} = \exp\left(i\sum_{n\geq 0} \tilde{a}_{I,n}(L_n + L_{-n} - \bar{L}_n - \bar{L}_{-n})\right) \quad (3.18)$$

with a_n, \tilde{a}_n real and L_n, \bar{L}_n representing the standard Virasoro generators. The notation a_n, \tilde{a}_n should not be taken to imply any relation between these coefficients. The form of a given V_a is universal in the sense that the coefficients $a_n, \tilde{a}_n, \hat{\beta}$ are fully determined by the size and shape of the a th hole h_a , and thus by the moduli of Σ , and do not depend on the particular choice of CFT. Thus for general L_a we may still write the state as in (3.12) with the V_a given by (3.17). Due to the U, \tilde{U} factors the V_a are not generally diagonal in the energy basis. But as above we find $U(h_a), \tilde{U}(h_a) \rightarrow \mathbf{1}$ and $\hat{\beta}_a \approx \beta_a + \beta_a^0$ in any limit where h_a shrinks to a point. Here the β_a^0 are determined by those moduli which are held constant.

3.2 More boundaries and factorization limits

The above discussion of the three-boundary state can easily be extended to general $n > 3$. The main novelty is that the surface Σ then has internal moduli, which correspond to conformally invariant cross ratios in the CFT correlators. In the puncture limit $L_a \rightarrow 0$, holding the internal moduli τ_α fixed, the treatment of the vanishing L_a at each puncture proceeds as above and we find

$$|\Sigma\rangle = \sum_{i_1 \dots i_n} C_{i_1 \dots i_n}(\tau_\alpha) e^{-\frac{1}{2}\sum_a \beta_a H_a} |i_1\rangle_1 \dots |i_n\rangle_n, \quad (3.19)$$

where the coefficients $C_{i_1 \dots i_n}$ are related to the CFT n -point functions. The dependence of $|\Sigma\rangle$ on the τ_α includes contributions from the dependence of the CFT n -point functions on the cross ratios, factors from the possible rotation in V_a , and the β_a^0 term in (3.3). We have absorbed both of these latter factors into the definition of $C_{i_1 \dots i_n}$ since they become independent of L_a in the puncture limit.

The most interesting aspect of the dependence on internal moduli is that the state exhibits factorization limits, which we now discuss in some detail. For simplicity, we restrict to the

four boundary case and corresponding factorization limits of the four-point function. The surface Σ is depicted in figure 3 and the internal moduli can be taken to be L_{14} and θ_{14} .

In the limit $L_{14} \rightarrow 0$ for fixed L_a (which need not be small), the surface Σ splits into two spheres, each with two holes, that remain connected by a thin tube. This corresponds to an s -channel limit of the corresponding four-point function. For a CFT with a unique ground state, at leading order in this limit only the vacuum propagates in this channel and the four-point function is dominated by the contribution from the disconnected two-point functions. The state then factorizes:

$$|\Sigma\rangle \approx |\Sigma_{14}\rangle|\Sigma_{23}\rangle, \quad (3.20)$$

where each factor is given by a path integral over an appropriate two-boundary manifold Σ_{14} or Σ_{23} . These two-boundary manifolds are constructed via a two-step procedure: First cut the original 4-boundary Σ along $L_{14} = L_{23}$ to obtain two disconnected pieces with three boundaries each. Now cap off the L_{14}, L_{23} boundary of each piece by sewing in a disk; i.e., by attaching the vacuum state at this boundary.

Each of the resulting surfaces Σ_{14}, Σ_{23} is then topologically an annulus. So as above we may write the state $|\Sigma_{14}\rangle$ in the form

$$|\Sigma_{14}\rangle = U_1 U_4 \left(\sum_i e^{-\frac{1}{2}\hat{\beta}_{14} E_i} |i\rangle_1 |i\rangle_4 \right) \quad (3.21)$$

for some $\hat{\beta}_{14}$ and some unitary operators U_1, U_4 of the form (3.18) acting respectively on boundaries 1 and 4. These operators and the parameter $\hat{\beta}_{14}$ in (3.21) are determined by the two remaining moduli L_1, L_4 . There is of course a corresponding expression for $|\Sigma_{23}\rangle$.

As in section 3.1, taking the puncture limit¹⁰ causes the length of each arm to diverge and the Weyl factor Ω governing the conformal transformation to become rotationally invariant near each boundary. As a result, the operators U_1, U_4 become trivial. Thus, in the puncture limit the state $|\Sigma_{14}\rangle$ reduces to a thermofield double state as one might expect. However, the temperature $\hat{\beta}_{14}$ of this state is more surprising. We may also follow section 3.1 in deducing the asymptotic form of $\hat{\beta}_{14}$. In the arm that includes each boundary, we may again use the BTZ approximation (3.8) for $\rho < \rho_0$ to conclude that in mapping Σ_{14} to a sphere with two round holes, the radii of the holes will be $r_1 = r_0 e^{-\frac{1}{2}\beta_1}, r_4 = r_0 e^{-\frac{1}{2}\beta_4}$. We therefore find

$$\hat{\beta}_{14} = \beta_1 + \beta_4 + \beta_0. \quad (3.22)$$

We will argue in section 3.3 that in fact $\beta_0 = 0$ and further that $\hat{\beta}_{14} = \beta_1 + \beta_4$ and even far from the puncture limit (at least when $\beta_1 = \beta_4$). With appropriate conventions (again at least when $\beta_1 = \beta_4$) it will turn out that $U_1 = U_4 = \mathbf{1}$ as well.

These moduli β_1, β_4 are specifying the geometry of Σ_{14} ; it has two exterior regions, which are $t = 0$ surfaces for BTZ geometries at the inverse temperatures β_1, β_4 . In the puncture

¹⁰Here we imagine taking $L_1, L_4 \rightarrow 0$ but maintaining $L_1, L_4 \gg L_{14}$.

limit, we find that the CFT path integral on Σ_{14} defines a thermofield double state at the inverse temperature $\hat{\beta}_{14}$. Taking $\beta_0 = 0$, this gives

$$T_{14} = T_1 T_4 / (T_1 + T_4), \quad (3.23)$$

where T_1, T_4 are the temperatures of the black holes in regions 1 and 4 of the bulk 4-boundary connected wormhole specified by our moduli. This relation between the temperatures is somewhat surprising. One might have guessed that at least for $\beta_1 = \beta_4$ the temperature of the thermofield double would agree with those in the connected 4-boundary wormhole. But this is not what we find.

It is interesting to compare this CFT analysis of factorization to the discussion of bulk phase transitions in section 2.3. In the bulk we expect a phase transition at some finite value of L_{14} : as we decrease L_{14} at fixed (sufficiently large) values of the other parameters, we expect to pass from the regime dominated by the Lorentz-signature four-boundary wormhole to one dominated by two disconnected two-boundary wormholes.¹¹ This bulk phase transition implies an approximate factorization of the state; the entropy of the reduced density matrix ρ_{14} on $CFT_1 \otimes CFT_4$ changes from being proportional to the central charge (on the side represented by connection in the bulk) to $\mathcal{O}(1)$ (on the side represented by a disconnected bulk). But since the two copies of BTZ are still connected through the Euclidean geometry, this factorization is only approximate. There is still some CFT entanglement, which is reflected in quantum entanglement between bulk fields on the disconnected Lorentzian geometries. It is only in the limit of vanishing L_{14} that the Euclidean separation between the two copies of BTZ diverges, and the state factorizes as seen above.

From this bulk perspective, we can interpret (3.23) as determining the temperature of the bulk BTZ geometry which provides the saddle-point for the bulk path integral. It again feels somewhat surprising that the resulting BTZ temperature is related to the original moduli in this nontrivial way. It would be very interesting to derive the temperature (3.23) directly from the bulk path integral, perhaps using the technology of e.g. [48, 49].¹² Applying the same reasoning to the 3-boundary phase transition of section 2.3 from the connected 3-boundary wormhole to a BTZ-plus-global-AdS₃ phase shows that the BTZ temperature is again related to the black hole temperatures of the connected wormhole by the analogue of (3.23).

In addition to the above s -channel limit, there are corresponding t - and u -channel limits. The t -channel limit $L_{12} \rightarrow 0$ corresponds to $L_{14} \rightarrow \infty$ with $\theta_{14} = 0$. In this limit the state will factorize as $|\Sigma\rangle = |\Sigma_{12}\rangle|\Sigma_{34}\rangle$. The u -channel limit $L_{13} \rightarrow 0$ corresponds to $L_{14} \rightarrow \infty$ with $\theta_{14} = \pi$. In this limit the state will factorize as $|\Sigma\rangle = |\Sigma_{13}\rangle|\Sigma_{24}\rangle$. In each case, the factorized states take the same form as in (3.21).

¹¹Smaller values of L_a may be dominated by four copies of global thermal AdS which pair up to form two thermofield double states, or perhaps by a single thermofield-double pair of global AdS₃ spaces and a single BTZ spacetime. The details depend on the choice of spin structures as in appendix A.2.

¹²See [41] for discussion of a related case in which the naive period of imaginary time also disagreed with the physical temperature of the Lorentz-signature black hole.

3.3 Reduced density matrices and factorization

To further explore the structure of our CFT states it is useful to consider the reduced density matrices obtained by taking partial traces over subsets of the degrees of freedom. In our multiboundary context, a simple class of reduced density matrices is obtained by tracing over one or more of the boundaries. These reduced density matrices also have a simple interpretation in terms of the path integral, and again exhibit interesting factorization properties.

For example, starting with the three-boundary state (3.4) and tracing over CFT_3 leads to the density matrix

$$\rho_{12} = \sum_{ij'j'} \rho_{ij'j'} |i\rangle_1 |j\rangle_2 \langle i'|_1 \langle j'|_2 \quad (3.24)$$

with $\rho_{ij'j'} = \sum_k A_{ijk} A_{i'j'k}^*$. As A_{ijk} is given by a CFT path integral over the pair of pants, this density matrix is a corresponding path integral over the 4-boundary space defined by sewing together two identical pairs of pants along the pair of boundaries corresponding to CFT_3 . Note that we do not integrate over the parameter L_3 associated with that boundary; this becomes a modulus for the resulting Riemann surface. The density matrix $\rho_{ij'j'}$ thus corresponds to a path integral over the sphere with four holes and so is related to a CFT four-point function. In contrast to the previous discussion of the four-point function this reduced density matrix still only depends on the three parameters labeling our original pair of pants geometry. Thus the reduced density matrix involves only a three-dimensional slice in the six-dimensional moduli space of Riemann surfaces with four boundaries.

Nevertheless, we can again take an s -channel limit; this corresponds to $L_3 \rightarrow 0$, so that the two pairs of pants factorize along the identification. Thus the density matrix (3.24) factorizes into the product of a bra-vector and a ket-vector and describes a pure state. From the point of view of the original state (3.4), this is because taking $L_3 \rightarrow 0$ sends the coefficients A_{ijk} to zero except when $|k\rangle_3$ is the vacuum state ($k = \mathbb{1}$). The nonzero coefficients $A_{ij\mathbb{1}}$ are then determined by the path integral over a manifold conformal to an annulus as in section 3.2. Indeed, (3.24) becomes

$$\rho_{12} \approx |\Sigma_{12}\rangle \langle \Sigma_{12}| \quad (3.25)$$

with $|\Sigma_{12}\rangle$ given by (3.21) (replacing 4 by 2). The corresponding puncture limit is described by (3.23) and $U_1, U_2 \rightarrow \mathbb{1}$ so that we then obtain a thermofield double state entangling boundaries 1 and 2. What is happening here is that the limit $L_3 \rightarrow 0$ forces the entropy of the accessible degrees of freedom in the third Hilbert space to be small. The other two spaces are then mostly entangled with each other and very little with the third.

We can again compare with the dual bulk description, where we expect that reducing L_3 for fixed L_1, L_2 should lead to a phase transition at some finite value of L_3 where we pass from the three-boundary black hole to a two-boundary wormhole and a copy of global AdS.

The other interesting limit for the density matrix is $L_3 \rightarrow \infty$. Since the twist is fixed to $\theta = 0$, this corresponds to the t -channel limit for the four-point function. This again gives a degeneration limit of the sphere with four holes, but where the boundaries are paired

differently than in (3.25). Thus in this limit the density matrix ρ_{12} factorizes into a density matrix for boundary 1 and a density matrix for boundary 2 with little correlation:¹³

$$\rho_{12} \approx \rho_1 \otimes \rho_2. \quad (3.26)$$

Reasoning as before, the density matrix ρ_1 will be given by a bulk path integral over a two-boundary Riemann surface whose geometry depends just on L_1 , which gives for general L_1 a result of the form

$$\rho_1 = U_1 e^{-\hat{\beta}_1(L_1)H_1} U_1^{-1}, \quad (3.27)$$

where U_1 has the form (3.18). There will be a similar result for ρ_2 . The factorization of ρ_{12} implies that the 3-boundary state for large L_3 takes the approximate form

$$|\Sigma\rangle \approx U_1 U_2 e^{-\frac{1}{2}\hat{\beta}_1 H_1} e^{-\frac{1}{2}\hat{\beta}_2 H_2} \sum_{ij} |i\rangle_1 |i\rangle_{31} |j\rangle_2 |j\rangle_{32} \otimes |\gamma\rangle_{33}, \quad (3.28)$$

where as usual $U_1, U_2 \rightarrow \mathbb{1}$ as $L_1, L_2 \rightarrow 0$. In writing (3.28) we have factored the CFT₃ Hilbert space \mathcal{H}_3 into a product of three Hilbert spaces $\mathcal{H}_3 \approx \mathcal{H}_{31} \otimes \mathcal{H}_{32} \otimes \mathcal{H}_{33}$ and introduced bases $\{|i\rangle_{31}\}, \{|j\rangle_{32}\}$ for $\mathcal{H}_{31}, \mathcal{H}_{32}$ (which are not energy eigenstates) as well as a state $|\gamma\rangle_{33} \in \mathcal{H}_{33}$. By leaving arbitrary the precise choices of these factor spaces and bases (as well as the choice of $|\gamma\rangle_{33}$) we have been able to remove the operator U_3 that one might have expected from (3.12) and (3.17).

The above t -channel factorization property is again easy to understand from elementary considerations in the overall Hilbert space. In this limit the space of accessible degrees of freedom in the third Hilbert space is much larger than for the other two, so in direct analogy with the classic results of Page [36], each of the remaining Hilbert spaces is strongly entangled with the large space. This leaves no room for significant entanglement between the small Hilbert spaces and the reduced density matrix ρ_{12} approximately factorizes into $\rho_1 \otimes \rho_2$.

Unlike the previous factorization limits, this limit is not associated with a phase transition in the bulk geometry; the connected 3-boundary wormhole remains the dominant bulk saddle even though the CFT state takes the approximate product form (3.28), where there is little entanglement between boundaries 1 and 2. We will see in section 4 that it is instead associated with a change in minimal surface in the calculation of the entanglement entropies from the bulk.

Thus, tripartite entanglement is not a necessary condition to connect three asymptotic regions in the bulk. Indeed, any notion of intrinsically tripartite entanglement will by definition vanish for (3.28) and we may say that the entanglement in this factorization limit is entirely bipartite. This expressly rules out any conjecture that GHZ-like states always play an essential role in the CFT description of connected multiboundary wormholes. Note that

¹³This particular factorization limit is special in that the above approximations are equally valid for all choices of spin structure. Since the shrinking geodesic separates two copies of boundary 1 from two copies of boundary 2, the sewing construction forces fermions on this cycle to have antiperiodic boundary conditions. This may be seen for example from the fact that the 4-boundary path integral would vanish if periodic boundary conditions were assigned. See appendix A.2 for further discussion.

the approximation in (3.28) relies on large L_3 but not large c . In particular, even corrections at order c^0 are exponentially suppressed at large L_3 .

We could now consider our usual analysis of the puncture limit to determine $\hat{\beta}_1$ for small L_1 . But in the present context we can do much better by recalling that the bulk solution is the connected 3-boundary wormhole described in the conformal frame that makes the boundary stress tensor rotationally symmetric – and thus time-independent – on each boundary, at leading order in the central charge c . However, other than a pure rotation, any non-trivial conformal transformation U_1 in (3.27) would add boundary gravitons, deforming the boundary stress tensor calculated in (3.27) away from spherical symmetry and adding non-trivial time-dependence. Since pure rotations commute with the Hamiltonian, we may therefore take $U_1 = \mathbb{1}$ in (3.27). By appropriate choices of phase for the states $|i\rangle_{31}, |j\rangle_{32}$ we may take $U_1 = \mathbb{1}, U_2 = \mathbb{1}$ in (3.28) as well. Noting that the expected stress tensor must match that of the bulk black hole also fixes $\hat{\beta}_1 = \beta_1$, and similarly $\hat{\beta}_2 = \beta_2$. These results hold for all L_1, L_2 and not just in the limits $L_1, L_2 \rightarrow 0$. Furthermore, since $U_1, U_2, \hat{\beta}_1, \hat{\beta}_2$ are determined entirely by conformal transformations independent of the central charge, these results are exact; they hold for any finite central charge as well as in the limit $c \rightarrow \infty$. As a result, while we began only with the assumption that stress tensor expectation values are thermal at leading order in large c , thermality of the the full reduced density matrix at this large L_3 limit is an exact result at all finite c .

This argument may also be used to determine the Boltzmann factor $\hat{\beta}_{14}$ and conformal transformations U_1, U_4 in the two-boundary states $|\Sigma_{14}\rangle, |\Sigma_{23}\rangle$ from (3.20)(3.21) that describe the approximate factorization of the 4-boundary state $|\Sigma_{1234}\rangle$ at small L_{14} . The key observation is that path integral computing the density matrix

$$\rho_{12} = \text{tr}_{\mathcal{H}_3} [|\Sigma_{123}\rangle\langle\Sigma_{123}|] \quad (3.29)$$

defined by tracing the 3-boundary state $|\Sigma_{123}\rangle$ over \mathcal{H}_3 coincides with that defining $|\Sigma_{1234}\rangle$. Here we make appropriate restrictions on the moduli of Σ_{1234} and reinterpret bra vectors on two of the four boundaries as kets. Explicitly, we may identify 1 and 4 in $|\Sigma_{1234}\rangle$ with the bra and ket copies of 1 in ρ_{12} , and 2 and 3 in $|\Sigma_{1234}\rangle$ with the bra and ket copies of 2 in ρ_{12} . Then the limit $L_3 \rightarrow \infty$ corresponds to the s -channel limit of section 3.2, and the path integral which gives ρ_1 in (3.27) coincides with the one for (3.21) when $\beta_1 = \beta_4$. Comparing (3.27) and (3.21) then requires

$$\hat{\beta}_{14} = 2\beta_1 = \beta_1 + \beta_4 \quad (3.30)$$

for any $\beta_1 = \beta_4$ and forces U_1, U_4 to be pure rotations. Thus $U_1 = \mathbb{1} = U_4$ for appropriate choices of the origin $\phi = 0$ of the angular coordinate on each boundary. In particular, this must occur when we choose $\phi = 0$ on boundaries 1 and 4 to represent points related by the \mathbb{Z}_2 reflection symmetry of Σ_{14} that exchanges boundaries 1 and 4 while mapping the L_{14} boundary to itself with two fixed points. It is tempting to speculate that (3.30) might hold even when β_1, β_4 are distinct.

For equal β_1, β_4 the relation (3.30) must in particular hold in the puncture limit where it must agree with (3.22). This implies that β_0 vanishes in (3.22), since β_0 is independent

of β_1, β_4 in this limit. It would clearly be of interest to check this result directly by finding the desired conformal transformation. That a consistent value of β_0 can be found at all constitutes a useful check on our arguments, and in particular verifies the absence of a factor of $1/2$ in (3.22), so that $\hat{\beta}_{14}$ is the sum of β_1 and β_4 rather than the average.

Returning to the discussion of tracing over boundaries, we note that reduced density matrices obtained by tracing over more than one boundary are related to CFT path integrals on surfaces of higher genus.¹⁴ For example, taking the three-boundary state (3.4) and tracing over two of the boundaries leads to the density matrix $\rho = \sum_{ii'} \rho_{ii'} |i\rangle\langle i'|$ with

$$\rho_{ii'} = \sum_{jk} A_{ijk} A_{i'jk}^*, \quad (3.31)$$

which corresponds to a path integral on the torus with two holes. This density matrix is a function of three moduli, so it corresponds to a three-dimensional subspace of the associated moduli space. The absence of Dehn twists implies that this is a slice of the moduli space where the torus is rectangular, so that the A and B cycles are orthogonal. In the puncture limit, this density matrix is related to the CFT two-point function on the torus.

4 Multipartite holographic entanglement

We saw in section 3 that the n -boundary state factorizes in various limits. Some limits are associated with phase transitions in the bulk geometry, so that factorization occurs in regions of moduli space where the dominant bulk saddle-point describes a disconnected Lorentz-signature spacetime. Notably, however, the $L_3 \rightarrow \infty$ limit of section 3.3 leads to a CFT state of the approximate product form (3.28), with correspondingly negligible entanglement between CFT_1 and CFT_2 , in a region where the connected 3-boundary wormhole remains the dominant bulk saddle.

This section will further probe the entanglement structure of the n -boundary states by using the Hubeny-Rangamani-Takayanagi (HRT) generalization [6] of the Ryu-Takayanagi proposal [5] to compute entropies holographically. We confine ourselves to regions of moduli space for which the dominant Lorentz-signature bulk phase is the connected n -boundary black hole, so these results are largely complementary to the preceding discussion of factorization limits. In some cases this may require imposing periodic boundary conditions on fermions as discussed in appendix A.2. Due to our emphasis on bulk methods below, this section will describe entropy as being associated with a given boundary B_a rather than a given CFT copy CFT_a as in section 3.

We will see that even when the geometry remains connected, there are phase transitions corresponding to changes in the HRT extremal surface. Such transitions indicate changes in the leading-order entanglement structure at large central charge. In particular, in the three-boundary case we find a phase transition associated with the large L_3 factorization

¹⁴We could also have encountered path integrals over higher genus surfaces simply by considering black holes with topology behind the horizon. We have avoided this further complication for simplicity of the discussion, though the topic will be addressed in [43].

limit, where the $\mathcal{O}(c)$ component of the mutual information between boundaries 1 and 2 vanishes for $L_3 \geq L_1 + L_2$. The relation of the phase transition to the factorization limit in the CFT is similar to that for the changes of bulk geometry: the phase transition occurs at finite L_3 , long before we get to the factorization limit described in section 3.3. Thus we learn that essentially bipartite entanglement (at leading order in the central charge) suffices to generate a 3-boundary connected wormhole over a sizeable part of the moduli space. Conversely, we will also find points in moduli space where the entanglement is fully tripartite. Intriguingly, at the level we are able to probe in this work we find that our entanglement results always coincide with properties of appropriate random states defined by sets of constraints associated with minimal geodesics.

For much of section 4, we will be somewhat cavalier in calling correlation as diagnosed by the mutual information and its tripartite analog “entanglement”. However, in quantum information theory, entanglement refers very specifically to correlation that cannot be generated without quantum mechanical interaction [50, 51]. While there is reason to believe that these measures do diagnose proper entanglement in holographic theories [23], no proof exists. Therefore, to solidify our claims, we repeat some of our analysis in section 4.4 quantifying entanglement more rigorously. Doing so leads to similar conclusions.

We begin by reviewing the holographic calculation of entanglement entropy in section 4.1. Sections 4.2 and 4.3 then study phase transitions between different minimal surfaces as a function of the bulk parameters and show that the bipartite vs. multipartite nature of the dual state depends strongly (and perhaps surprisingly) on the moduli of Σ . We also use these calculations to fix the parameters $\hat{\beta}_1, \hat{\beta}_2$ in (3.28), as well as β_0 in (3.22). Section 4.4 addresses correlation vs. entanglement, and section 4.5 discusses k -party entanglement in the n -boundary state for larger k, n and shows that the *distillable* entanglement between the asymptotic regions is at least $(n+1)/2$ -partite. Section 4.6 develops new tools to characterize n -partite entanglement and uses these tools to show that for wormholes with n similar sized horizons there is n -partite entanglement for even n and $n-1$ -partite entanglement for odd n . Section 4.7 then notes that, so far as the above probes can tell, the nature of the entanglement of the state is consistent with what one would expect for suitably generic states. We also discuss extensions of this random-state model beyond the regime investigated thus far.

4.1 Holographic calculation of entanglement entropy

For time-independent states invariant under time-reversal, the entropy of spatial subregions may be studied holographically using the Ryu-Takayanagi proposal [5]. This prescription states that, given a spatial subregion A in a constant-time surface of the CFT, the associated reduced density matrix ρ_A obtained by tracing over degrees of freedom in the complement of A has an entropy S_A given by the area (i.e., the co-dimension two notion of volume, so in fact a length on AdS_3) of a minimal surface γ_A in the associated bulk constant-time surface whose boundary coincides with that of A ($\partial\gamma_A = \partial A$) and for which γ_A is homologous to A . The symmetries imply that the bulk is static, so there is a unique bulk surface orthogonal to the time-translation which contains A . This proposal has been extensively tested and explored in the literature, and a strong connection to Euclidean quantum gravity arguments

was made in [7].

This construction must be generalized for time-dependent states such as our (3.1). The HRT prescription [6] replaces the minimal surface γ_A in the spatial slice with an extremal surface \mathcal{E}_A in the bulk spacetime which is homologous to A . As evidence in its favor, this generalization has been shown to satisfy strong subadditivity [52].

In addition, while it may be difficult to extend the argument of [7] to the case of fully general time-dependence, the extension to subregions of the time-reflection invariant $t = 0$ surface in states like the ones we consider appears straightforward. Since each state (3.1) is defined by a Euclidean path integral, replica partition functions may be defined in the usual way for any subregion A of the $t = 0$ slice. Following the rest of the argument of [7], and in particular using the assumption that replica symmetry is unbroken, then leads one to consider conical singularities in the bulk $t = 0$ surface. Since our system is invariant under time-reversal, the Euclidean bulk $t = 0$ surface also appears in the bulk Lorentzian geometry. Since this surface is positive definite, any homology class contains a minimal surface. And since time-reversal invariance sets the extrinsic curvature of our $t = 0$ surface to zero, this minimal surface is an HRT extremal surface. So at least for regions A lying at $t = 0$, we arrive at the HRT proposal so long as there are no further extremal surfaces with smaller area¹⁵. In principle one should be able to move A away from $t = 0$ by an appropriate analytic continuation of the Euclidean results, though we will have no need to pursue this complication below; studying the behavior of sub-region entropy under Lorentzian evolution is an interesting problem for the future.

We shall restrict attention here to the case where A consists of one or more of the boundaries in their entirety. Unitarity then implies that the entanglement entropy is independent of time and the above argument applies directly. Thus we compute the entropy $S(B_1 \cup \dots \cup B_k)$ of the reduced density matrix on some collection of asymptotic boundaries. For brevity we write this entropy as simply $S(B_1 \dots B_k)$. Since A is a closed manifold, the corresponding bulk minimal surface γ_A must be a collection of closed geodesics in Σ which together are homologous to $B_1 \cup \dots \cup B_k$. One natural candidate is the union $H_1 \cup \dots \cup H_k$ of the horizons associated with the boundaries over which we do not trace, though there will also be others involving internal geodesics and other horizons. One alternative that always exists is the union $H_{k+1} \cup \dots \cup H_n$ of the horizons in the regions over which we trace. Since the overall state (3.1) is pure, we always have $S(B_1 \dots B_k) = S(B_{k+1} \dots B_n)$, and $S(B_1 \dots B_n) = 0$.

From this basic data we can construct quantities such as the mutual information,

$$I(A : B) = S(A) + S(B) - S(AB) \tag{4.1}$$

the triple information,

$$\begin{aligned} I_3(A : B : C) &= I(A : B) + I(A : C) - I(A : BC) \\ &= S(A) + S(B) + S(C) - S(AB) - S(BC) - S(AC) + S(ABC) \end{aligned} \tag{4.2}$$

¹⁵While we have no sharp argument to forbid such surfaces, we note that any one such surface must spontaneously break the time-reversal symmetry. In addition, under the assumptions of [52] they must be spacelike separated from the $t = 0$ extremal surfaces and also, in an appropriate sense, farther from the boundary. It seems unlikely that such surfaces exist in our geometries. We will henceforth assume that they do not.

and higher-order generalizations. The mutual information is nonnegative and equal to zero if and only if the state factorizes. I_3 , on the other hand, can take either sign in general, but is always nonnegative in holographic theories [23]. As such, its magnitude measures how much additional information there is in BC about A that is not present in B and C separately.

Note that in the subsequent discussion we will find that in some regimes of parameters, various mutual informations vanish. What this really means is that at the leading order in central charge probed by the HRT formula the mutual information is vanishing; we do not expect the value to be precisely zero, but rather some value of order one, which is not visible in the holographic analysis.

4.2 Phases of entanglement for three boundaries

As we vary the parameters of Σ , there can be transitions where the relative area of two candidate minimal surfaces changes sign, so that control over the HRT entanglement passes from one to the other. These phase transitions reflect changes in the entanglement structure of the dual CFT state at leading order in the central charge.

Consider the three-boundary case, with L_i parametrizing the size of the horizon of region i . As described above, the HRT minimal surfaces that compute entanglement are unions of these horizons. In this case, there is a transition at $L_3 = L_1 + L_2$; if $L_1 + L_2 < L_3$, the minimal surface for $B_1 \cup B_2$ and B_3 is $H_1 \cup H_2$, so

$$S(B_1 B_2) = S(B_3) = \frac{\pi}{2G}(L_1 + L_2). \quad (4.3)$$

But for $L_1 + L_2 > L_3$, the minimal surface for $B_1 \cup B_2$ and B_3 is H_3 , so

$$S(B_1 B_2) = S(B_3) = \frac{\pi}{2G}L_3. \quad (4.4)$$

It is natural to associate this phase transition with the factorization limit $L_3 \rightarrow \infty$, since for $L_1 + L_2 < L_3$ the holographic mutual information satisfies

$$I(B_1 : B_2) = S(B_1) + S(B_2) - S(B_1 B_2) = \frac{\pi}{2G}(L_1 + L_2 - (L_1 + L_2)) = 0. \quad (4.5)$$

Thus there is no entanglement of boundaries 1 and 2 at leading order in central charge whenever they are small subsystems, not just in the strict decoupling limit. Taking $L_1 \leq L_2$, there is also a further phase transition when $L_2 > L_1 + L_3$, where the minimal surface for B_2 changes from H_2 to $H_1 \cup H_3$. Thus, setting $L_{tot} = L_1 + L_2 + L_3$ and choosing $L_1 \leq L_2$, the mutual information is

$$I(B_1 : B_2) = \frac{\pi}{2G} \times \begin{cases} 0 & \text{if } 2(L_1 + L_2) \leq L_{tot} \\ 2(L_1 + L_2) - L_{tot} & \text{if } 2(L_1 + L_2) > L_{tot} \text{ and } 2L_2 \leq L_{tot} \\ 2L_1 & \text{if } 2L_2 > L_{tot}. \end{cases} \quad (4.6)$$

In particular, the mutual information vanishes (at leading order in large central charge) until the horizons for black holes 1 and 2 are sufficiently large and then rises with slope 2 as a function of L_2 until saturating at $2L_1$.

When $L_3 > L_1 + L_2$, the entanglement seems almost entirely bipartite, between 1 and 3 and between 2 and 3, though we will argue this more carefully in section 4.4. In the extreme limit where $L_3 \rightarrow \infty$, we saw in the earlier field theory analysis that the entanglement becomes purely bipartite, and the reduced density matrix on tracing over 3 factorizes. Here we see that there is a sizeable region in the parameter space where bipartite entanglement dominates. This strengthens the conclusion of section 3.3: a large multipartite entanglement component is not required to generate a branched wormhole connecting multiple asymptotic regions.

Of course, the relation to the CFT three-point function suggests that the entanglement for finite L_3 is not purely bipartite; we expect a subdominant tripartite component. But this can be compared to the situation in different geometric phases: When the bulk saddle is three copies of AdS, the holographic entanglement entropy calculation of course says that the mutual information between any two copies vanishes. The actual CFT state will have some order one entanglement between the two copies, but this subleading entanglement does not appear to generate a geometric connection. Similarly, here the subdominant contribution presumably plays no important role in generating the connected geometry. Thus it appears to be enough that B_1 and B_2 have large entanglement with different parts of the B_3 Hilbert space for the bulk dual to have a geometric connection.

On the other hand, when $L_3 < L_1 + L_2$, there seems to be more than purely bipartite entanglement. This is an example of a general phenomenon; the extent to which our CFT state involves multipartite entanglement depends on the moduli of Σ . To discuss the multipartite nature of the entanglement in detail, we will turn below to cases with more boundaries, where this can be more clearly diagnosed. But first we make some further remarks on this three-boundary example.

We note that the leading-order behaviour (4.6) matches the mutual information for typical pure states chosen according to the unitarily invariant measure on $\overline{\mathcal{H}}_1 \otimes \overline{\mathcal{H}}_2 \otimes \overline{\mathcal{H}}_3$, where $\overline{\mathcal{H}}_a$ is a finite-dimensional Hilbert space of dimension $e^{\pi L_a/2G}$. This may be seen by using the results of [36] to compute the entropies $S(B_1)$, $S(B_2)$, $S(B_1 B_2)$ and assembling these building blocks to find (4.6). Thus, for the three-boundary case, the nature of the entanglement between subsystems is what we would expect for a “random” entangled state. Let us stress that this is not to say that the state *is* generic; the state dual to the three-boundary wormhole is very special. It is just the very coarse overall entanglement structure that appears to be generic. We will discuss this issue further in section 4.7.

It is also interesting to note that the fact that $S(B_3) = \frac{\pi}{2G}(L_1 + L_2)$ implies that knowledge of the reduced density matrix ρ_3 is telling us about the geometry beyond the horizon H_3 for this asymptotic region. Indeed, in general the HRT surface corresponding to a region A in the boundary lies outside the causal wedge associated to A in the bulk spacetime [52]. These black holes provide a particularly striking example of this behaviour. In the BTZ case, the entropy of the reduced density matrix on one side is related to the area of the horizon. Here, it is related either to the area of the horizon H_3 , or to the sum of the areas of the horizons H_1 and H_2 .¹⁶

¹⁶Note that the minimal surface will always lie somewhere in the shadow region, or perhaps on its boundary. This places the surface in the part of the space that is causally inaccessible from any of the boundaries and

4.3 Phases of entanglement for four boundaries

We now explore the phases of holographic entanglement in the four boundary case, for which the triple mutual information provides a useful probe of truly multipartite correlations.

Let us first consider the limit of small L_a , where we can assume that the length of any internal geodesic is larger than that of any horizon. Thus, as in the three-boundary case, the only relevant minimal surfaces are those built from the horizons. We again consider only the connected 4-boundary wormhole in the bulk¹⁷.

There are then a number of different regions in the parameter space of such four-boundary black holes, distinguished by the behaviour of the mutual informations and triple informations. For example, if $L_1 + L_2 < L_3 + L_4$, the smallest minimal surface for $S(B_1 B_2) = S(B_3 B_4)$ will be $H_1 \cup H_2$, and the mutual information between boundaries 1 and 2 will vanish:

$$I(B_1 : B_2) = S(B_1) + S(B_2) - S(B_1 B_2) = 0 \quad (4.7)$$

as in the three-boundary case. Alternatively, if $L_1 + L_2 > L_3 + L_4$, the smallest minimal surface for $S(B_1 B_2) = S(B_3 B_4)$ will be $H_3 \cup H_4$, and then it is the mutual information $I(B_3 : B_4)$ which vanishes. If $L_1 + L_2 = L_3 + L_4$, both these mutual informations will vanish. There are three different such divisions of the set of boundaries into pairs, so there will be sixteen different regions of the parameter space, labeled by the different mutual informations which are nonzero. When all of the lengths are equal, $L_1 = L_2 = L_3 = L_4 = L$, all of the bipartite mutual informations vanish. This suggests that the entanglement in this case does not involve bipartite entanglement between the pairs of boundaries, but is truly multipartite. Indeed, we find

$$I_3(B_2 : B_3 : B_4) = I(B_2 : B_3) + I(B_2 : B_4) - I(B_2 : B_3 B_4) = -2L \cdot \frac{\pi}{2G} < 0, \quad (4.8)$$

which is in magnitude twice the entropy of an individual boundary, so I_3 is the most negative it can possibly be. (By strong subadditivity, $I_3(B_a : B_b : B_c) \geq -2S(B_a)$.)

There is also an interesting division of the parameter space associated with the triple informations. If any one of the lengths is greater than the sum of the others, all of the triple informations vanish. To see this, consider without loss of generality $L_1 > L_2 + L_3 + L_4$. Then the minimal surface for $B_a \cup B_b$ will be $H_a \cup H_b$ for $a, b = 2, 3, 4$, and the minimal surface for $B_2 \cup B_3 \cup B_4$ will be $H_2 \cup H_3 \cup H_4$. Thus all mutual informations between boundaries 2, 3 and 4 vanish: in particular $I(B_2 : B_3) = 0$, $I(B_2 : B_4) = 0$, and $I(B_2 : B_3 B_4) = S(B_2) + S(B_3 B_4) - S(B_2 B_3 B_4) = 0$, so $I_3(B_2 : B_3 : B_4) = 0$.

corresponds to the fact that the amount of entanglement cannot be affected by local quantum operations on any of the entangled systems. Acting with local operators on the boundaries can change the geometry in the regions that are causally accessible from the boundary, so it is natural that the minimal surface that calculates the entanglement entropy must lie outside these influenceable regions [52, 42].

¹⁷We expect that – at least by making appropriate choices of spin structure (see appendix A.2) and tuning moduli – one can arrange for this bulk phase to dominate the path integrals of section 2.2 somewhere in each of the regimes discussed below. But this remains to be shown in detail. If not, our analysis applies to whatever CFT state is in fact dual to these bulk spacetimes.

Moreover, although the mutual informations involving B_1 are nonzero, the triple mutual informations vanish: for example,

$$I(B_1 : B_2 : B_3) = S(B_1) + S(B_2) + S(B_3) - S(B_1B_2) - S(B_1B_3) - S(B_2B_3) + S(B_1B_2B_3), \quad (4.9)$$

but $S(B_1) = S(B_2) + S(B_3) + S(B_4)$, $S(B_1B_2) = S(B_3) + S(B_4)$, and $S(B_1B_3) = S(B_2) + S(B_4)$, and $S(B_1B_2B_3) = S(B_4)$, while $S(B_2B_3) = S(B_2) + S(B_3)$, so

$$\begin{aligned} I(B_1 : B_2 : B_3) &= S(B_2) + S(B_3) + S(B_4) + S(B_2) + S(B_3) \\ &\quad - S(B_3) - S(B_4) - S(B_2) - S(B_4) \\ &\quad - S(B_2) - S(B_3) + S(B_4) = 0. \end{aligned} \quad (4.10)$$

Similarly for $I(B_1 : B_2 : B_4)$ and $I(B_1 : B_3 : B_4)$. Thus, in this extreme part of the parameter space, all the triple mutual informations vanish. There are four such regimes, which are subsets of four of the sixteen regimes identified in the previous analysis of the mutual informations.

Thus, in the parameter space with four boundaries:

- There is a special point where all pairwise mutual informations vanish (always only to leading order) but at which the triple information is as negative as it can be. At that point, there is very little correlation between pairs of subsystems but there is correlation between any collection of three subsystems.
- There are regions where all the triple informations vanish, but that nonetheless have a large pairwise mutual information, which we interpret as indicating that the entanglement is essentially bipartite.

As a final example, again set all the lengths L_a to be equal to a fixed L and reintroduce L_{14} . As $L_{14} \rightarrow 0$, all other structures in figure 3 are pushed to the boundary, where the conformal factor diverges. Holding fixed the L_a then forces both L_{12} and L_{13} to become large, ensuring that of the internal minimal surfaces, only L_{14} can ever be small enough to be relevant in any holographic entropy calculation. One then finds

$$S(B_a) = \frac{\pi}{2G}L \quad (4.11)$$

$$S(B_1B_4) = S(B_2B_3) = \frac{\pi}{2G}L_{14} \quad (4.12)$$

$$S(B_aB_b) = \frac{\pi}{2G}2L \quad \text{for all } (a, b) \notin \{(1, 4), (2, 3)\} \quad (4.13)$$

$$S(B_aB_bB_c) = \frac{\pi}{2G}L \quad (4.14)$$

so that

$$I(B_1 : B_4) = I(B_2 : B_3) = \frac{\pi}{2G}(2L - L_{14}) \quad \text{as calculated earlier, and} \quad (4.15)$$

$$I(B_a : B_b) = 0 \quad \text{for all } (a, b) \notin \{(1, 4), (2, 3)\} \quad (4.16)$$

$$I_3(B_a : B_b : B_c) = -\frac{\pi}{2G}L_{14}. \quad (4.17)$$

We see that when $L_{14} = 2L$, there is no mutual information between any pair of boundaries and the triple informations are equal to $-\frac{\pi}{2G}2L$. Shrinking L_{14} linearly interpolates toward the other extreme, where there is only pairwise mutual information: In the factorization limit $L_{14} \rightarrow 0$, the triple informations vanish, leaving only pairwise maximal mutual information between B_1 and B_4 and between B_2 and B_3 . We may forbid a bulk phase transition in this limit to a pair of two-sided BTZ black holes by taking opposite spin structures for B_1 and B_4 (and thus also opposite spin structures for B_2 and B_3) as discussed in appendix A.2.

4.4 Entanglement versus classical correlation

The mutual information I quantifies bipartite correlations and, in holographic theories at least, the triple information I_3 provides a way of quantifying the extent to which correlations are intrinsically tripartite. None of these measures, however, distinguishes between correlation and entanglement. For pure states, there is no difference between the two, but a nontrivial distinction arises for mixed states. For example, consider the state $(|\uparrow\uparrow\rangle\langle\uparrow\uparrow| + |\downarrow\downarrow\rangle\langle\downarrow\downarrow|)/2$. While correlated, it is the statistical mixture of two product states so should not be considered entangled. Formally, a state is said to be entangled only if it cannot be decomposed as a convex combination of product states [50, 51]. In the holographic context, it is commonly assumed that bulk connectivity is related to true entanglement. However, we know of no definitive analysis of the role that might be played by mere correlation. We thus take care to distinguish the two below.

To begin, let us return to the 3-boundary black hole with $L_1 + L_2 < L_3$ for which we found that to leading order, $I(B_1 : B_2) = 0$, $I(B_1 : B_3) = 2S(B_1)$ and $I(B_2 : B_3) = 2S(B_2)$. From this, we provisionally drew the conclusion that the entanglement was essentially all bipartite between the pairs (B_1, B_3) and (B_2, B_3) . The extreme version of this situation is the $L_3 \rightarrow \infty$ limit discussed in section 2.3 in which the $B_1 B_2$ density operator factorizes. In that case, B_1 and B_2 are indeed each entangled only separately with B_3 . When L_3 is only slightly larger than $L_1 + L_2$, however, we expect that $I(B_1 : B_2)$ will be $\mathcal{O}(1)$ so the state will not factorize in general.

In what quantitative sense, then, can we say that most of the entanglement is bipartite? A stringent operational definition of entanglement is to ask for the maximum rate at which near-perfect Bell pairs can be extracted from many copies of a given state using only local operations and classical communication (LOCC), a quantity known as the distillable entanglement, E_D [51]. When a bipartite state is pure, E_D and the entanglement entropy are one and the same: $E_D(A : B) = S(A)$. For mixed states, however, the entanglement entropy is generally just an upper bound: $E_D(A : B) \leq \min\{S(A), S(B)\}$. In the case under consideration, $S(B_1) = \frac{\pi}{2G}L_1$ but there is very little mutual information $I(B_1 : B_2)$ so we expect $E_D(B_1 : B_2)$ to be small. A famous coding theorem in quantum information theory, on the other hand, gives the following “hashing” bound [53]:

$$E_D(A : B) \geq S(A) - S(AB). \quad (4.18)$$

Applying this inequality to the three boundary black hole gives

$$E_D(B_1 : B_3) \geq S(B_3) - S(B_1 B_3) \quad (4.19)$$

$$= S(B_1 B_2) - S(B_2) \quad (4.20)$$

$$= S(B_1) - I(B_1 : B_2). \quad (4.21)$$

So, when $L_1 + L_2 < L_3$, we get $E_D(B_1 : B_3) \geq S(B_1) - \mathcal{O}(1)$ and, similarly $E_D(B_2 : B_3) \geq S(B_2) - \mathcal{O}(1)$, confirming that B_1 and B_2 are each nearly as entangled as they can be with B_3 . (In fact, there is a single LOCC distillation procedure that will simultaneously extract both the $B_1 : B_3$ and the $B_2 : B_3$ entanglement at the specified rates in this case [54].)

So far, the distillable entanglement has simply confirmed the more heuristic mutual information-based analysis. However, states can be highly correlated and even entangled without being distillable. Applying the lower bound again, we find that in general, for $L_1 \leq L_2$,

$$\begin{aligned} E_D(B_1 : B_2) &\geq \max\{0, S(B_2) - S(B_1 B_2)\} \\ &= \frac{\pi}{2G} \times \begin{cases} 0 & \text{if } L_2 \leq L_3 \\ L_1 + 2L_2 - L_{tot} & \text{if } L_2 > L_3 \text{ and } 2L_2 \leq L_{tot} \\ L_1 & \text{if } 2L_2 > L_{tot}, \end{cases} \end{aligned} \quad (4.22)$$

where $L_{tot} = L_1 + L_2 + L_3$. Like the mutual information, this lower bound on E_D is piecewise linear as a function of L_2 . But it starts its increase from 0 at a later point, when $L_2 \geq L_3$ (or equivalently $L_1 + 2L_2 = L_{tot}$), nonetheless saturating at its maximum value $\frac{\pi}{2G} L_1$ when $2L_2 = L_{tot}$, just as does the mutual information. When E_D is positive, we can conclude unequivocally that there is bipartite entanglement between B_1 and B_2 . Otherwise, we can't be sure. Thus we cannot exclude the possibility that B_1 remains unentangled with B_2 in the regime close to, but on the $I(B_1 : B_2) > 0$ side of, the HRT phase transition.

In the analysis above, for reasons of convenience and conservatism, we have focused on the distillable entanglement. It is the smallest of the many inequivalent mixed state entanglement measures [55]. Another measure that might arguably be more relevant to holography is the entanglement of formation E_F [51], which is related to the minimal rate of Bell pairs required to produce near-perfect copies of the state in question using only LOCC [56]. In general, $E_F \geq E_D$ and the gap can be very large: the entanglement of formation can be near maximal even when the mutual information is very small [57]. So in the example above, it is possible that $E_F(B_1 : B_2)$ remains large even after the HRT phase transition leads to vanishing $I(B_1 : B_2)$. Since $E_F(B_1 : B_2)$ is defined as the minimum average entanglement entropy of the pure states in any convex decomposition of the mixed state ρ_{12} , this would be an indication that it is impossible to describe the connected 3-boundary Lorentzian wormhole without the use of entanglement even when the correlations between the B_1 and B_2 asymptotic regions are very weak. Unfortunately, HRT calculations alone can't be used to verify if that is the case because getting good lower bounds on E_F requires more detailed information about the structure of the state.

As in section 4.3, we close by considering the 4-boundary case with all horizon lengths equal to a fixed L and parameters chosen so that of the internal minimal surfaces only L_{14}

is small enough to play a role in HRT calculations. Using equations (4.11)-(4.14), it is easy to see that

$$E_D(B_1 : B_3 B_4) = \frac{\pi}{2G} L, \quad E_D(B_1 : B_3) = 0 \quad \text{and} \quad E_D(B_1 : B_4) \geq \frac{\pi}{2G} (L - L_{14}), \quad (4.23)$$

with equality holding in the first case because E_D can never exceed entanglement entropy and in the second because it can't exceed half the mutual information [58]. So the distillable entanglement between B_1 and $B_3 B_4$ is maximal, with the length of the internal minimal surface L_{14} controlling how much of the entanglement can be distilled using B_1 and B_4 alone. This is consistent with what we expect in the factorization limit $L_{14} \rightarrow 0$, but indicates that the distillable entanglement between B_1 and B_4 is significant even above any phase transition where the bulk geometry disconnects into a pair of two-sided BTZ black holes.

4.5 Entanglement structure as a function of n

Although the entanglement structure becomes more complicated as n increases, similar ideas can be used to develop a sense of the relative importance of bipartite, tripartite and higher order correlations in any given situation. Consider a black hole with n boundaries B_a and horizon lengths L_a sufficiently small that we can ignore internal minimal surfaces.

Since I_3 was a useful diagnostic for the 4-boundary case, it is tempting to evaluate its higher order generalization

$$I_k(B_1 : B_2 : \dots : B_k) = \sum_{\sigma} (-1)^{\sigma} S(\sigma), \quad (4.24)$$

where the sum is over subsets of the arguments $\{B_1, B_2, \dots, B_k\}$. Interpreting these functions as measures of multipartite correlation only makes sense, however, when each I_k is either always nonnegative or always nonpositive. An earlier computer search revealed that I_4 can be made to take both positive and negative values in holographic theories by choosing appropriate combinations of intervals in a single CFT [23]. One might hope that such complications do not occur in this case since each B_a is an entire CFT. But in fact it is even easier to find configurations in which I_4 takes either sign. When $L_a = L$ for all a , one finds that $I_4 = \frac{\pi}{2G} L$ for $n = 5$ but $I_4 = -\frac{\pi}{2G} 2L$ for $n = 6$.

With the naive approach ruled out, let us proceed instead using the more trustworthy mutual information and distillable entanglement functions. Set $L_{tot} = \sum_{a=1}^n L_a$ to be the total length of all the horizons and for convenience write $L_Z = \sum_{z \in Z} L_z$ for the sum of the horizons indexed by set Z . If $X, Y \subseteq \{1, \dots, n\}$ such that $L_X \leq L_Y$ then

$$I(B_X : B_Y) = \frac{\pi}{2G} \times \begin{cases} 0 & \text{if } 2(L_X + L_Y) \leq L_{tot} \\ 2(L_X + L_Y) - L_{tot} & \text{if } 2(L_X + L_Y) > L_{tot} \text{ and } 2L_Y \leq L_{tot} \\ 2L_X & \text{if } 2L_Y > L_{tot}. \end{cases} \quad (4.25)$$

This equation has exactly the same form as (4.6) and reflects the fact that the entanglement structure is similar to that of a generic random state. Likewise, we can consider the distillable

entanglement between collections of subsystems. Performing the same substitution gives

$$\begin{aligned}
 E_D(B_X : B_Y) &\geq \max\{0, S(B_Y) - S(B_X B_Y)\} \\
 &= \frac{\pi}{2G} \times \begin{cases} 0 & \text{if } L_X + 2L_Y \leq L_{tot} \\ L_X + 2L_Y - L_{tot} & \text{if } L_X + 2L_Y > L_{tot} \text{ and } 2L_Y \leq L_{tot} \\ L_X & \text{if } 2L_Y > L_{tot}. \end{cases}
 \end{aligned} \tag{4.26}$$

Combining these bounds on I and E_D does permit rigorous conclusions about multipartite entanglement.

For example, set $L_a = L$ for all a for simplicity. Then the bounds demonstrate that when $|X| + |Y| \leq n/2$, the mutual information between B_X and B_Y is small. Since $E_D \leq I/2$ [58], this means that the distillable entanglement of any subsystem of size $n/2$ or smaller is negligible. Conversely, if $|X| + 2|Y| > n$, then there *will* be distillable entanglement. In particular, substituting $|X| = 1$ tells us that this will be the case for any subsystem of size $|X| + |Y|$ strictly larger than $\frac{n+1}{2}$. In this precise sense, the distillable entanglement of the state is at least $\frac{n+1}{2}$ -partite: there is none for subsystems of size smaller than the threshold and there always is for subsystems larger.

Earlier, we used the triple information I_3 to infer that, in a 4-boundary black hole with all horizon lengths equal, all correlation was at least tripartite. But that argument was really at the level of correlation rather than entanglement. Nonetheless, the conclusion is consistent with our analysis here, which indicates that all the distillable entanglement is at least $\lceil \frac{4+1}{2} \rceil = 3$ -partite.

Finally, we note that the fact that the entropy of a pair of boundaries exceeds that of a single boundary means that GHZ-like n -party entanglement is unimportant.¹⁸

4.6 Intrinsically n -partite entanglement

Earlier, we identified the large- L_3 limit of a three-boundary black hole as an example of a system whose entanglement was not intrinsically tripartite because the state could secretly be factorized into a pair of bipartite entangled states, one between B_1 and B_3 and the other between B_2 and B_3 . In this section we will generalize that idea to define intrinsically k -partite entanglement and then show that some n -boundary black holes have definitely have some n -partite entanglement. If they didn't, then some of their HRT entropies would have to be significantly smaller than they are. However, we warn the reader that the argument below is not strong enough to quantitatively compare this n -party entanglement to k -party entanglement with $k < n$.

Our starting point will be that, for sufficiently small moduli L_x (with the index labeling boundaries now $x = 1, \dots, n$), the entropies calculated using the HRT formula resemble those of random quantum states in the sense that

$$S(\cup_{x \in X} B_x) = \min \left[\sum_{x \in X} S(B_x), \sum_{x \notin X} S(B_x) \right] \tag{4.27}$$

¹⁸For $n = 4$ one can also see this from (4.8), whose sign is opposite to that of the GHZ_4 state as is always the case for holographic I_3 [23].

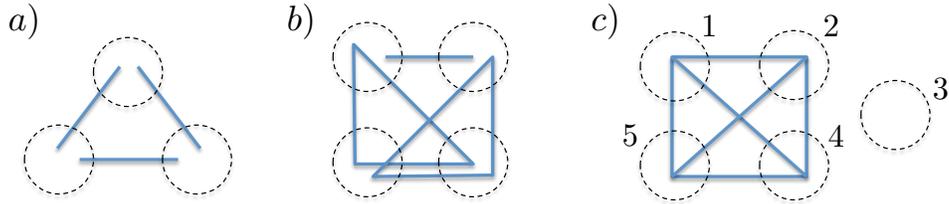


Figure 6: a) A 2-producible state on $B_1B_2B_3$. Each subsystem is indicated by a dashed circle, with component entangled states drawn as graphs, in this case lines, between the systems entangled by the state. Observe that if each 2-party state is a Bell state, then the overall tripartite state will have generic entropy behaviour. b) A 3-producible state composed of pure states entangled between pairs and triples of systems. c) A 4-producible state. If the state on $B_1B_2B_4B_5$ is generic, then the version of this state symmetrized over all 5 systems will have generic entanglement entropies.

to leading order in the central charge. We will again focus for simplicity on the case in which all the $S(B_x)$ are the same, so that the right hand side the formula is simply $\min(|X|, n - |X|)S(B_1)$. From now on, we will say that states satisfying (4.27) are *entanglement-generic*. Our objective will be to use (4.27) alone to demonstrate the existence of n -partite entanglement when n is even, and $(n - 1)$ -partite entanglement when n is odd.

To get a sense of the difficulties inherent in drawing conclusions from boundary entanglement entropies alone, consider figure 6a, which depicts a tripartite quantum state which is 2-producible, meaning that it is the tensor product of bipartite entangled states [59]. Unlike the $L_3 \rightarrow \infty$ example considered above, no 2-party reduced density matrix constructed from 6a will factorize. But among all tripartite states, the 2-producible ones are still very special and lack intrinsically tripartite entanglement. Nonetheless, choosing each bipartite factor to be maximally entangled is sufficient to make the states entanglement-generic. (The impossibility of detecting intrinsically tripartite entanglement in a pure tripartite state using entropies alone is reflected in the fact that the triple information I_3 will always be zero for such states.) Also, it is worth keeping in mind that, while 2-producible states contain only pairwise entanglement, starting with enough pairwise entanglement, one could make any tripartite state using only local operations and classical communication. The strategy would just be to manufacture the desired state entirely in one factor and then teleport the appropriate pieces to the other two. Figure 6b depicts a 3-producible state: it is the tensor product of states which are themselves entangled between at most three subsystems. If each of the three factors is chosen to be a pure state chosen at random from the Haar measure, the resulting state will not be entanglement-generic this time: the entropy of the top two subsystems will be strictly less than the sum by an amount equal to twice the entanglement entropy of the bipartite factor.

Figure 6c provides a final illuminating example. As drawn, it looks highly atypical: system B_3 factorizes from the rest. But suppose the $B_1B_2B_4B_5$ portion is selected according to the Haar measure and that all five B_x have the same dimension. Then the resulting state will with high probability be entanglement-generic. Because $S(B_3)$ will be zero, it doesn't

cause any problems; the entropy of any two out of the five subsystems will be the sum of their individual entropies. While the example treats B_3 very specially, it is straightforward to symmetrize it such that all $S(B_x)$ will be the same. Simply take a tensor product of five similar states, each one singling out a different B_x for factorization. The crucial fact that makes this trick work is that in our context (4.27) implies that the entropy of any X is determined by a collection involving at most half the subsystems. But five is an odd number, so that “at most half the subsystems” actually means two out of five and a Haar random state on four factors is generic on any two of its factors. As such, analogous constructions of entanglement-generic $(n-1)$ -producible states exist whenever n is odd. Therefore, we will at best be able to show that odd- n entanglement-generic systems cannot be $(n-2)$ -producible, meaning they are intrinsically $(n-1)$ -entangled.

Now consider the same construction in even dimensions. The entropy of a fixed subset in the symmetrized model is the same as the entropy of a random subset in the unsymmetrized model. So suppose that $B_1 B_2 \cdots B_{n-1}$ is in a random pure state factorizing with the state of B_n . A random subset of $B_1 B_2 \cdots B_n$ of size $n/2$ will contain B_n with probability $1/2$. If it does, the entropy, instead of being $nS(B_1)/2$, will be $(n/2-1)S(B_1)$. Otherwise, the entropy is as it should be for an entanglement-generic state. The expected deficit will therefore be $S(B_1)/2$. In the symmetrized model, the entropy of a system of size $n/2$ will therefore be $S(B_1)/2$ smaller than would be required to be entanglement generic. Our analysis below will show that this is optimal to within a constant factor: every $(n-1)$ -producible state will exhibit an average entropy deficit on subsystems consisting of $n/2$ factors of at least an n -independent constant times $\frac{1}{n} \sum_x S(B_x)$.

Before proceeding, let us introduce some definitions and notation in the interest of being more precise. Let $[n] = \{1, \dots, n\}$ and \mathcal{C}_j^X consist of the j -element subsets of X . In the rest of the section, B_x will refer specifically to a Hilbert space. A state $|\psi\rangle \in \otimes_{x=1}^n B_x$ is said to be k -producible [59] if there exist Hilbert spaces B_x^K and isometries $U_x : B_x \rightarrow \otimes_{K \in \mathcal{C}_k^{[n]}} B_x^K$ such that

$$\otimes_{x \in [n]} U_x |\psi\rangle = \otimes_{K \in \mathcal{C}_k^{[n]}} |\psi^K\rangle \quad (4.28)$$

for some $|\psi^K\rangle \in \otimes_{x \in K} B_x^K$.¹⁹ While this definition is transparent and completely sufficient, it will be very convenient below to be able to refer to B_x^K for $x \notin K$ with the understanding that B_x^K is trivial in that case. So we will instead work with the equivalent definition that $|\psi^K\rangle \in \otimes_{x \in [n]} B_x^K$, with $B_x^K = \mathbb{C}$ for $x \notin K$.

For any Hilbert spaces B_x , write $B_X = \otimes_{x \in X} B_x$. Our goal will be to quantify deviations from being entanglement-generic, which means we will be comparing $S(B_X)$ to $S_{\text{sum}}(B_X) = \sum_{x \in X} S(B_x)$. Our main result will be a lower bound on the difference $S_{\text{sum}}(B_X) - S(B_X)$ in terms of $|X|$. The main idea will be to generalize the observation made above about n -partite states composed of generic $(n-1)$ -producible factors. The challenge is to prove a lower bound that works *only* assuming that the state is k -producible and nothing else about its internal structure. (As noted above, however, we assume in this section that the entropies $S(B_x)$ are identical for all subsystems $x \in [n]$. The bound holds regardless but

¹⁹This definition is a slight generalization of the usual notion, which doesn't explicitly allow for the isometries U_x .

that assumption simplifies its interpretation.)

The derivation of the bound is quite technical so we defer the argument to appendix B, with the precise statement appearing as (B.41). The result may be expressed in terms of a fractional entropy deficit $\delta(j, k, n)$:

$$\overline{[S_{\text{sum}}(B_X)]}_{X \in \mathcal{C}_j^{[n]}} - \overline{[S(B_X)]}_{X \in \mathcal{C}_j^{[n]}} \geq \delta(j, k, n) \overline{[S(B_x)]}_{x \in [n]}, \quad (4.29)$$

where the notation $\overline{[\cdot]}_{X \in \mathcal{C}_j^{[n]}}$ denotes the average over j -element subsets. This average is of course trivial for $S_{\text{sum}}(X)$, as we have assumed that all $X \in \mathcal{C}_j^{[n]}$ give the same $S_{\text{sum}}(X)$. Equation (B.41) shows that $\delta(j, k, n)$ is usually positive for allowed j, k, n but does not give immediate insight into its magnitude so we have evaluated it on a computer.

Consider first $(n - 1)$ -producible states when n is even. The strongest test is when $j = n/2$. As shown in figure 7, δ appears to be bounded below by a constant independent of n : there will always be an appreciable fractional entropy deficit. The smallest deficit, $2/9$, occurs when $n = 4$, with the values up to $n = 100$ being as high as 0.138 and converging to a value in the vicinity of $1/8$. Since the symmetrized $(n - 1)$ -producible state constructed in the introduction to this section had an entropy deficit of $1/2$, the lower bound is within a constant factor of being optimal.

When n is odd, the fractional entropy deficit for $(n - 1)$ -producible states is always 0, reflecting our earlier observation that it is possible to construct $(n - 1)$ -producible entanglement-generic states. There is an appreciable fractional entropy deficit for $k = n - 2$, however. For $n = 5$, one finds a value of $1/20$, for example, and it appears to converge to a value in the vicinity of 0.03.

Studying the bound as a function of k is unfortunately not as satisfying. Since being k -producible becomes more and more stringent as k decreases, one would expect the bound to likewise get stronger. This happens up to a point, as is visible in figure 7. The bound generally strengthens in going $k = n - 1$ (or $k = n - 2$ when n is odd) to $k = n - 2$ ($k = n - 3$ when n is odd) but thereafter gets weaker as k decreases. That isn't a surprise since the bound was designed to demonstrate the existence of intrinsically n -partite entanglement. Nonetheless, it would be interesting to improve the analysis so as to extract good bounds for all k .

In summary, we have confirmed that when all L_x are equal, and when the HRT surfaces are unions of horizons, the state $|\Sigma\rangle$ is n -partite entangled for even n and $(n - 1)$ -partite entangled for odd n . If it weren't then the average entanglement entropies evaluated using the HRT formula of subsystems of size $n/2$ would be smaller by an amount proportional to the central charge and independent of n than their actual values. Our assumption about the HRT surfaces will certainly hold in the picture limit.

4.7 The random state model

Our CFT states have an interesting entanglement structure that depends on the choice of moduli. While the details may seem complicated, we noted in section 4.2 that the leading-order bipartite mutual information (4.6) for the 3-boundary wormhole is precisely that found

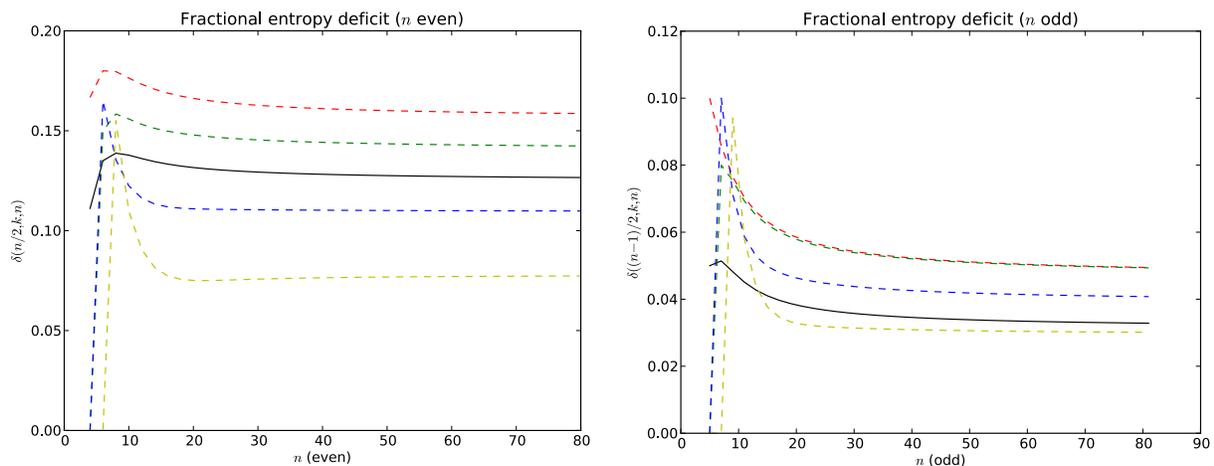


Figure 7: Lower bounds on the fractional entropy deficits as a function of n . The plot on the left depicts $\delta(n/2, k, n)$ for n even. The solid black line corresponds to $k = n - 1$ and the dashed red, green, blue and yellow lines to $k = n - 2$, $k = n - 3$, $k = n - 4$ and $k = n - 5$, respectively. In all cases, δ converges to a constant, with the strongest bound actually occurring for $k = n - 2$. The plot on the right depicts $\delta((n - 1)/2, k, n)$ for n odd. In this case, the deficit for $k = n - 1$ is precisely zero so the solid black line corresponds instead to $k = n - 2$. The dashed red, green, blue and yellow lines in turn represent $k = n - 3$ through $k = n - 6$. Once again, the deficits converge to constants at large n . By definition, the actual entropy deficits can only increase as k decreases so these plots illustrate that our bounds could be strengthened for small k : the dashed yellow line, which in both plots corresponds to the smallest value of k , is also in both plots the weakest bound.

in the large L_a/G limit for a random state on the Hilbert space $\overline{\mathcal{H}}_1 \otimes \overline{\mathcal{H}}_2 \otimes \overline{\mathcal{H}}_3$, where $\overline{\mathcal{H}}_a$ is a finite-dimensional Hilbert space of dimension $e^{\pi L_a/2G}$. Such Hilbert spaces provide natural finite-dimensional models of our CFTs in the connected 3-boundary wormhole phase where the leading-order entropy of each black hole is just $\pi L_a/2G$. For this reason we may think of the states in $\overline{\mathcal{H}}_a$ as all having the same energy, namely that corresponding to the a^{th} black hole. In the puncture limit we might read off this conclusion directly from the Boltzmann factors in the CFT wavefunction.

In just the same way we may again use [36] to see that our general mutual information formula (4.25) for an arbitrary number of boundaries – and thus also the triple mutual information (4.8) for arbitrary L_a – agree with corresponding random state models at large L_a/G so long as the internal geodesics can be ignored. The point is again that any subsystem with less than half the total dimension has nearly maximal entropy, while purity of the full state requires the entropy of larger subsystems to be equal to that of their complement.

This result seems rather natural for a strongly interacting CFT with a high density of states. It indicates that the n -point functions have rather little structure when viewed at a sufficiently coarse-grained level. Of course, the actual n -point functions in any given CFT are definite objects and, in the same way, the n -boundary states are rather specific states. This is all in direct analogy with the two-boundary case where the resulting thermofield-double state might be approximated by the fully entangled microcanonical ensemble state

$$|\psi\rangle_{micro} = \sum_{E \in (E_1, E_2)} |E\rangle|E\rangle \quad (4.30)$$

defined by an appropriate range of energies. The entropy of each factor, and indeed the full reduced density matrix on each factor, agrees to leading order with that of a random state chosen with respect to the unitarily invariant measure on the appropriate $\overline{\mathcal{H}}_1 \otimes \overline{\mathcal{H}}_2$. Nonetheless it is clear that the actual state (4.30) requires very specific correlations between the energies of the two systems that will not be reproduced by generic states. See e.g. [38, 60, 61] for further discussion of this point.

Taking this model seriously makes predictions for both the distillable entanglement E_D and the entanglement of formation E_F . The most striking of these governs the behaviour of E_F in the three-boundary case. When $L_1 + L_2 > L_3$, we have seen that $I(B_1 : B_2) \propto L_1 + L_2 - L_3$ so that the correlation between B_1 and B_2 shrinks as L_3 increases. In the associated random state model, however, E_F does not shrink, instead staying essentially constant and maximal right until $L_1 + L_2 = L_3$, at which point it drops abruptly to a negligible level [57]. If the same is true for the entanglement of formation of the CFT states, then as long as $I(B_1 : B_2)$ is of order the central charge, any description of the state ρ_{12} as a mixture of pure states would involve only nearly maximally entangled states. This would be particularly interesting if those maximally entangled states could be interpreted holographically as different connected two-boundary spacetimes. The consequences for E_D of the random state model are less striking, just that the lower bounds calculated in sections 4.4 and 4.5 are tight modulo violations of the so-called additivity conjecture [57], which are believed to be small for the random state model [62].

One would very much like to test these predictions – or indeed those associated with

any other entanglement measure – via further bulk calculations. It would be particularly interesting to study cases with a large number of boundaries n . While the entanglement properties predicted by the simplest random state model will be invariant under arbitrary permutations of the boundaries, any choice of Σ will break this symmetry. Working in the puncture limit we may note that placing any n points on the sphere to define an n -point function will induce a notion of proximity, so that some subsets of points may be said to be closer together than others. One would very much like to understand the extent to which this manifests itself in entanglement properties of the CFT state and leads to discrepancies with the random state model.

Now, as we have already seen, even at small n the state becomes somewhat less generic in regions of moduli space where the internal geodesics become relevant. These internal geodesics effectively function as constraints that restrict the structure of the state. In the presence of such constraints one might construct a natural random state model by first decomposing the relevant n -boundary manifold Σ into pairs of pants by cutting Σ along some set of internal geodesics associated with the moduli L_{ab} . Considered separately, the state on each pair of pants would then be drawn randomly from the associated Haar-random distribution as above. However, there are several issues to consider further:

- i) There is in general some correlation between states on the I th and J th pair of pants. For example, if the moduli of these surfaces agree, the states are identical. It would be interesting to understand the extent to which such correlations are important for reproducing various properties of the n -boundary state. It would also be of great interest to understand how rapidly such correlations decay as the moduli are deformed away from equality. One should similarly study how these correlations depend on the Dehn twist parameters θ along the cuts.
- ii) The decomposition of Σ into pairs of pants is not unique. In simple cases, in a given region of moduli space, it may be possible to single out the decomposition associated with the tightest possible set of constraints. But in general we expect different decompositions to restrict the state in different ways. For a given prescription for dealing with point **i** above, the A th decomposition effectively defines some measure μ_A on the space of n -boundary states. So we must combine these μ_A into a single measure μ_Σ . While there is no canonical way to do so, one may hope that some universality renders the details unimportant for large L_a . One simple recipe that is then to write each measure as $\mu_A = f_A \mu_{Haar}$ in terms of the Haar measure on n boundaries and, noting that any Σ has a finite number of decompositions A , to define $\mu_\Sigma = \mathcal{N} (\prod_A f_A) \mu_{Haar}$ by simply multiplying the functions f_A and choosing the coefficient \mathcal{N} to properly normalize the result. As desired, this μ_Σ vanishes on any state forbidden by any pair of pants decomposition.
- iii) When the decomposition involves many pairs of pants, one might begin to question the microcanonical approximation in which the Boltzmann factors of (3.17) are replaced by truncation to a finite-dimensional Hilbert space. One might like to instead explore

properties of states chosen randomly from a Gibbs ensemble with some specified density of states.

A class of such models for certain domains of moduli space will be explored in the follow-up paper [63].

5 Conclusions

Our work has focused on the holographic description of CFT states $|\Sigma\rangle$ defined by the Euclidean path integral over Riemann surfaces Σ with n asymptotic boundaries. Such states may be interpreted as entangled states of n separate CFTs each defined on the cylinder $S^1 \times \mathbb{R}$; i.e., as states of $\text{CFT}_1 \otimes \cdots \otimes \text{CFT}_n$. This study was motivated by the fact that – at least in appropriate regions of moduli space – the dual bulk states describe time-symmetric black hole geometries where the $t = 0$ surface is precisely the surface Σ . Our main goal was to understand the relation of these geometries to entanglement in $|\Sigma\rangle$ as a function of both external and internal moduli (L_a and τ_α). We found a number of interesting features.

First, we noted in section 2.3 that the multiboundary black hole is not always the dominant saddle-point in the bulk semi-classical approximation. As the moduli of Σ vary, there can be phase transitions between this maximally connected phase and phases where only certain subsets of boundaries are connected in the bulk. As with any first-order phase transition, these are associated with discontinuities in the entropy and energy of each asymptotic region (and thus of each CFT factor). But by studying a factorization limit, we found in section 3.2 that any notion of temperature one might assign is also discontinuous across the phase boundary. A particular striking example occurs when a pair of boundaries splits off from the rest but remain connected to each other. The two-boundary factor is then (conformal to) a thermofield double state, but with inverse temperature given by the sum of those for the corresponding black holes in the maximally connected phase. The disconnection of an $m > 2$ boundary component – perhaps the remainder of the geometry from which the above pair of boundaries was severed – will also feature a discontinuity in all temperatures, though a milder one that vanishes in the puncture limit of small moduli L_a . When a single boundary disconnects, it is unclear whether a meaningful notion of temperature can be associated with the resulting factor at all.

While unfamiliar, this phenomenon should be expected since our path integrals do not generally describe thermal equilibrium. There is thus no canonical assignment of temperatures to a given point in moduli space. Instead, any approximate notion of temperature must emerge from the physics in the same sense that we may expect generic pure states to equilibrate to thermal states at late times in sufficiently ergodic systems with many degrees of freedom. This places the notion of temperature in our context on a footing similar to those of energy and entropy so that such discontinuities are natural. Our understanding of the above phase transitions was based on CFT calculations and indirect arguments from an HRT calculation; it would be very interesting to understand the structure better in the bulk.

Another interesting result was the strong moduli-dependence of the multipartite nature of $|\Sigma\rangle$. Consider for example the three-boundary case. In the limit $L_3 \rightarrow \infty$ we found

CFT₁ and CFT₂ to be essentially unentangled – and indeed, uncorrelated; see section 4.4 for further discussion of correlation vs. entanglement. Holographically we found that the mutual information between 1 and 2 vanished at leading order in the central charge whenever $L_3 > L_1 + L_2$. Thus the separate entanglement of systems 1 and 2 with system 3 appears sufficient to produce a geometric connection between 1 and 2; multiparty entanglement is unimportant in this regime.

One might have expected this bipartite entanglement to produce two separate wormholes, one linking 1 to 3 and one linking 2 to 3, but in AdS₃ it is not possible to have two separate black hole horizons in a single asymptotic region; the two must lie inside a single horizon [64]. In higher dimensions one expects²⁰ to be able to construct states having only separate 1,3 and 2,3 bipartite entanglement for which the holographic dual does indeed have two initially-separate wormholes which link 3 to 1 and 2 respectively. But at least in the absence of angular momentum, gravitational attraction will eventually cause the two black holes to coalesce and form a connected multiboundary wormhole with properties similar to those considered here. Whether or not this possibility is realized in a given state may depend on the spatial separation of the degrees of freedom in 3 that are entangled with 1 and 2. Similar extreme regions of moduli space with purely bipartite entanglement (where $n - 1$ smaller systems are each entangled only with one big system at leading order in c) exist for any number of boundaries n . In contrast, for $n = 4$ or more boundaries there are complementary regions of moduli space where all bipartite mutual informations between pairs of boundaries vanish.

For typical parameter values the degree of multipartite entanglement appears to increase with the number of boundaries. In particular, for all horizon lengths equal (and when the internal moduli impose no constraints) the entanglement includes at least some n -partite entanglement when n is even and $(n - 1)$ -partite entanglement when n is odd. Thus, multipartite entanglement is increasingly important, but the dominant form of entanglement may involve only a subset of the boundaries connected geometrically in the bulk. The techniques developed here to reach those conclusions, namely bounding the entanglement entropy of k -producible quantum states, may be of independent interest.

Intriguingly, all of the above entanglement results in both bipartite and multipartite domains agree precisely with the predictions of a simple random-state model. As discussed in section 4.7, extensions of this model to cases with important internal constraints are more subtle. These cases remain to be explored in detail, though some initial results will appear in [63].

Our studies were facilitated by writing the state $|\Sigma\rangle$ as a sum of products of CFT n -point functions of operators O_i and the states made by the action of certain operators V_a , one for each boundary, on the state made by the action of O_i on the vacuum (see Sec. 3.1). The V_a are universal, in the sense that they are (complexified) conformal transformations dictated only by the moduli of Σ and independent of any details of the CFT. Each V_a in fact takes the form $U_a e^{-\frac{1}{2}\hat{\beta}_a H_a} \tilde{U}_a$, where U_a, \tilde{U}_a are unitary conformal transformations acting on CFT_a and H_a is the associated Hamiltonian. Although the moduli-dependence of U_a, \tilde{U}_a and the

²⁰Though see comments in section 3 of [52].

inverse temperature $\hat{\beta}_a$ may be complicated in general, it simplifies in the puncture limit $L_a \rightarrow 0$ where one finds $U_a, \tilde{U}_a \rightarrow \mathbb{1}$ and – up to a constant offset – $\hat{\beta}$ is just the inverse temperature of the corresponding bulk black hole in the fully-connected single-wormhole bulk geometry. Thus we obtain an explicit formula in this limit.

We also remind the reader of the interesting phase structure we found in the 3-boundary case when we increased L_3 with L_1, L_2 held fixed. Due to the transition of the minimal surface discussed in section 4.2, for $L_3 \geq L_1 + L_2$ the mutual information $I(B_1 : B_2)$ is only order 1 (as opposed to $\mathcal{O}(c)$). Thus for most probes of CFT_1 and CFT_2 the reduced density matrix ρ_{12} obtained by tracing out B_3 is well approximated by the product $\rho_1 \otimes \rho_2$. But as noted in section 3.3, the bulk geometry in this region is still the three-boundary wormhole. Although particularly dramatic here, this structure is similar to that associated with other Ryu-Takayanagi transitions where one finds related-but-distinct transitions for various Renyi entropies S_n though the bulk geometry for the state changes continuously [48].

Despite the above results, many features of $|\Sigma\rangle$ remain to be understood. Interesting avenues for further work include developing better measures of multipartite entanglement – which could diagnose more clearly the extent of multipartiteness in our states – and using them to explore the moduli dependence of the entanglement in more detail. Even restricting attention to simple measures, it seems likely that one can make further use of HRT to explore the structure of (3.28) by sewing two identical pairs of pants together along boundary 3 with a general Dehn twist θ .

One would also like to gain some further insight into details of the CFT states. For example, one might study model systems for which the n -point functions are explicitly known. In another direction, numerical computations may be of use in understanding both the conformal transformations U, \tilde{U} and the inverse temperature $\hat{\beta}_a$ in (3.17) beyond the puncture limit. More generally, it would be desirable to obtain a precise characterization of those CFT states whose bulk duals are described by a single bulk geometry. It may be that such states exhibit specific information theoretic features – such as the monogamy of mutual informations – that distinguish them from generic CFT states.

Acknowledgements

We thank David Berenstein, Steve Giddings, Tom Hartman, Sean Hartnoll, Juan Maldacena, John Parker, Joe Polchinski, Grant Salton, Volkher Scholz, Kostas Skenderis, Lenny Susskind, Nate Thomas, and David Tong for useful conversations. V.B. was supported in part by DOE grant DE-FG02-05ER-41367. V.B. also thanks the Fondation Pierre-Gilles de Gennes for support, and enjoyed the hospitality of the LPT at the Ecole Normale Supérieure and the LPTHE at the Université Pierre et Marie Curie in the early stages of this project. P.H. was supported in part by CIFAR, FQXi, NSERC and the Templeton foundation. A.M. was supported in part by NSERC, the Simons Foundation, FQXi and the Templeton foundation. D.M. was supported in part by the National Science Foundation under Grant No PHY11-25915, by FQXi grant FRP3-1338, by funds from the University of California, and as a Visiting Fellow Commoner at Trinity College during much of this work. In addition, he

would like to thank DAMTP, Cambridge U. for their hospitality throughout much of this work. S.F.R. is supported by STFC. D.M. and S.F.R. also thank the Isaac Newton Institute for its hospitality during critical stages of the project.

A Details on Bulk Phases and the Mapping Class Group

A.1 Mapping Class Group

The phase structure of section 2.3 can be understood a bit more formally as follows. We will begin with the torus case (for which see [65, 45, 66] for details). The Euclidean geometry is, topologically, a solid donut whose boundary is a torus. This torus has two non-contractible cycles – the ϕ circle and the Euclidean time t_E circle – which we will call a and b , respectively, with intersection number $a \cap b = 1$. One can then consider the action of the mapping class group (MCG) – the group of diffeomorphisms, modulo those which are continuously connected to the trivial diffeomorphism – on the torus. This mapping class group can be most easily visualized by considering its action on $H^1(T^2)$, i.e. on the cycles a and b of the boundary torus. It must act linearly on these cycles, and it must preserve the intersection number $a \cap b = 1$. Therefore an element of the MCG can be regarded as an element of $SP(2, \mathbb{Z}) \cong SL(2, \mathbb{Z})$. This is the usual modular group of the torus.

The action of the MCG on the bulk solid donut is a diffeomorphism, which reduces to the former diffeomorphism on the boundary torus. As this diffeomorphism is non-trivial at the boundary, it is not a gauge transformation of the bulk theory, and the action of the MCG is non-trivial in the bulk; these should be thought of as distinct bulk saddles. For example, it takes the geometry where the a cycle is contractible in the bulk (thermal AdS) to a geometry where a linear combination of a and b are contractible. Thus we end up with a family of geometries labelled by the elements of $SL(2, \mathbb{Z})$. Typically, one of these will dominate, but at certain points in moduli space two of these saddles might have equal action. Since saddles are related by the action of the MCG , this will happen only when the boundary torus is a fixed point of the MCG . At this point the boundary torus has a nontrivial discrete symmetry (an automorphism), since it is invariant under the action of an element of the MCG . The inverse-temperature β can be interpreted as the imaginary part of the conformal structure modulus of the torus, via $\text{Im } \tau = 2\pi\beta$. The modulus τ is acted on by $SL(2, \mathbb{Z})$ in the usual way, and has a fixed point at $\tau = i$. This is precisely the Hawking-Page transition. It is important to note that the two different Euclidean saddle points – thermal AdS and the Euclidean BTZ black hole – have the same metric, but that they are related by a diffeomorphism which acts nontrivially on the boundary, so should be regarded as giving independent contributions to the path integral of the theory.

For Riemann surfaces of higher genus, the story is similar. For a genus g Riemann surface M , the non-contractible cycles in $H^1(M)$ are similarly divided into a cycles and b cycles, a_i and b^i , with intersections $a_i \cap b^j = \delta_i^j$. The action of the MCG on the cycles is determined by the modular group $Sp(2g, \mathbb{Z})$, which acts linearly on the a and b cycles and preserves the intersection number. For higher genus, some elements of the MCG act trivially on the a and b cycles; mathematically, this is the statement that the Torelli group (the quotient of

the mapping class group by the modular group) is nontrivial. However, our interest is on the action on the cycles, as this is what affects the interpretation of the bulk handlebody.

In the bulk handlebody, half of these cycles will become contractible. There is a choice of division where the cycles that are contractible in the bulk are all a cycles. The action of the MCG will then map this to different combinations of a and b cycles being contractible in the bulk. Again, this is a diffeomorphism of the bulk handlebody, but these should be thought of as distinct bulk saddles, and will have different physical interpretations.

Focusing on the case of genus two, the modular group is $Sp(4, \mathbb{Z})$. There are four fundamental cycles, a_1, a_2, b_1, b_2 . Choices for which cycles become contractible in the bulk correspond to the different phases discussed in section 2.3 and shown in Figure 4. The cases of interest, which have moments of time reversal symmetry, are:

- (a_1, a_2) contractible. This is the naive handlebody obtained by filling in the genus 2 Riemann surface when we embed it in \mathbb{R}^3 . This is depicted on the left hand side of Figure 4. This analytically continues in Lorentzian signature to three disconnected AdS's.
- (b_1, b_2) contractible. This is the "dual" of the naive handlebody described above; it is the exterior of the Riemann surface, if we think of the Riemann surface as embedded in \mathbb{R}^3 along with a point at infinity. This is depicted in the middle of Figure 4. This geometry analytically continues to the connected wormhole.
- $(a_1, b_2), (a_2, b_1)$ or $(a_1 + a_2, b_1 + b_2)$ contractible. These three cases correspond to three different AdS + BTZ geometries; the different choices determine which pair of horizons is connected by a BTZ wormhole. The case (a_1, b_2) contractible is shown on the right hand side of Figure 4.

The different bulk saddles are related by the action of $Sp(4, \mathbb{Z})$, and the phase transitions where bulk saddles exchange dominance must occur at fixed points of $Sp(4, \mathbb{Z})$. In appendix A.3 below, we use this to identify the location of the Hawking-Page like phase transition between the naive handlebody and its dual in a particularly simple subspace of the moduli space of genus two surfaces.

A.2 Spin structures

Section 2.3 discussed the phase structure without worrying about boundary conditions for fermions. But explicit realizations of the AdS₃/CFT₂ duality contain fermions in both the CFT and the dual bulk system. These fermions require a choice of boundary conditions. Each copy of the CFT lives on a circle on which we can have either periodic or antiperiodic boundary conditions corresponding to the Ramond (R) or Neveu-Schwarz (NS) sectors.

This choice of boundary conditions makes an important difference to the field theory dynamics. In particular, as we now review, defining a CFT state on n cylinders by a path integral over Σ leads to a vanishing result unless the spin structures chosen on the n circles satisfy certain constraints. From the bulk point of view, the choice of boundary conditions

influences which bulk saddles can contribute as they must admit a consistent spin structure for bulk fermions.

Thinking of Σ as a sphere with n holes shows that the discussion is equivalent to that associated with computing n -point functions. The conformal transformation that maps the CFT from the cylinder to the plane flips the fermion periodicity, so a boundary with periodic (R sector) boundary conditions corresponds to the insertion of an R sector operator, which has the property that fermionic operators on the worldsheet pick up a minus sign on transport around a circle enclosing the R sector operator. Thus, an R sector operator has an associated branch cut for fermions. Since a branch cut on the sphere must have two ends, the nonzero n -point functions must involve an even number of R sector operators. In the puncture limit this is a familiar statement, but it is clearly true more generally: the CFT path integral defining our state is nonzero only if we have an even number of boundaries with R boundary conditions. So, for example, in the two-boundary case we can either have NS on both boundaries or R on both boundaries. This is not surprising; two boundaries connected by a cylinder must have the same behavior. For the three-boundary case, the possible fermion boundary conditions are NS on all three boundaries, or NS on one boundary and R on the others.

For the CFTs of interest in string theory, the NS sector has a unique vacuum state, so the behaviour in the limits $L_3 \rightarrow 0, \infty$ is as discussed in the main text. But the behavior in the R case is more subtle. There is a degenerate space of R vacuum states, so if we choose a periodic spin structure on the degenerating cycle, the R vacuum states would propagate along this channel, and the density matrix would not exactly factorize as above. However, the space of R vacuum states is parametrically smaller at large c than the space of states at any finite temperature, so the picture remains qualitatively similar: In the $L_3 \rightarrow 0$ limit the entropy of ρ would drop from that of a macroscopic black hole to the smaller entropy obtained by counting the R ground states.

Another important effect, however, is that the logarithmic density of states (measured by the entropy) near the R vacuum is of order c times a coefficient that vanishes at precisely zero energy. Thus the limiting $L_3 \rightarrow 0$ behavior becomes visible only at extremely small moduli with $L_3/\ell \lesssim e^{-Ac}$ for some A of order 1. Since the bulk semi-classical approximation breaks down in this regime, holography provides no tools for studying this limit.²¹ In particular, there is no reason to expect semi-classical bulk physics to exhibit an associated phase transition.

Indeed, the choice of fermion boundary conditions influences which of the possible bulk saddle-points discussed in section 2.3 can contribute, as the bulk saddle must admit fermion fields consistent with the choice of boundary conditions on each boundary. In differential geometry, the choice of boundary conditions for fermion fields on a non-contractible cycle in spacetime is referred to as a choice of spin structure. Since fermions must have antiperiodic boundary conditions on a contractible cycle, consistency requires fermions to pick up a minus sign under a 2π rotation.

For the two-boundary case, the possible bulk saddles were BTZ or two copies of global

²¹The problem is closely related to distinguishing black hole microstates using bulk methods.

AdS₃. BTZ has spin structures consistent with either boundary condition (both NS or both R), but in double-global-AdS₃ saddle both spatial circles shrink to zero in the interior and so require the NS boundary condition. We see that the Hawking-Page like phase transition between these geometries occurs only for the both-NS case. For the R boundary conditions the BTZ black hole dominates at all temperatures [45].

For the three-boundary case, imposing NS boundary conditions on all boundaries similarly allows all of the bulk saddles discussed in section 2.3. But for one NS and two R boundaries we are not allowed to fill in the R boundaries with global AdS₃. Thus there are only two possible saddles given by the three-boundary connected wormhole and the BTZ-plus-global-AdS₃ saddle with the AdS₃ factor having the NS boundary. There should be a single phase transition between them as one varies the length of the geodesic associated with the NS boundary.

The story is similar for $n > 3$ boundaries: the connected n -boundary black hole is an allowed bulk saddle for all consistent choices of boundary conditions, and the remaining saddles describe disconnected Lorentzian bulk geometries with the i th component connecting m_i boundaries of which an even number have R boundary conditions. So long as Σ has genus zero, the choice of fermion boundary conditions on the asymptotic circles in fact fully determines the spin structure on Σ . This can be seen from the fact that any minimal geodesic cuts Σ into two disconnected pieces. Since the total number of R asymptotic boundaries is even, both pieces will have the same number of R asymptotic boundaries counted modulo 2. But we obtain a nonzero path integral on each piece only if the total number of R boundaries on each side is even. So an even number of R asymptotic boundaries on one side requires an NS cut, and an odd number requires an R cut. This affects the factorization limits of section 3.2 in direct parallel with the discussion of phase transitions above, with factorization associated with degenerating R cycles being invisible to semi-classical bulk physics.

However, we note that there is a special case immune from such concerns. This arose in our discussion of the ρ_{12} density matrix obtained from the 3-boundary state by tracing over boundary 3 in the limit $L_3 \rightarrow \infty$. There the shrinking cycle cut Σ into two pieces for which *each* piece was separately invariant under a \mathbb{Z}_2 symmetry that exchanged boundaries in pairs. Thus each piece always has an even number of R boundaries, and the cut must be NS. It follows that there is always an associated phase transition at some finite-but-large value of L_3 .

A.3 Hawking-Page like phase transition for three boundaries

Section 2.3 argued that varying the moduli L_a of our 3-boundary state should lead to certain phase transitions in the dual bulk geometries. The simplest transition to analyze occurs when all moduli are equal ($L_1 = L_2 = L_3 = L$) and involves the exchange of relative dominance of the 3-global-AdS₃ phase and the connected 3-boundary wormhole phase. If the BTZ + global AdS₃ phase were to dominate, it would constitute spontaneous breaking of the S_3 permutation symmetries on such boundaries. This seems unlikely to occur, and we will neglect this phase below.

From the Euclidean point of view, this phase transition can be viewed as a change in the

preferred way of filling in the genus two surface, which is the doubled version of our path integral over the sphere with three holes. In the two-boundary case, the analogous surface was a torus, and we know that the transition took place at a point of enhanced symmetry, where the torus was square, with both cycles of the same length. Therefore here rather than trying to explicitly calculate the Euclidean action for the two different saddles to determine the transition point, we will simply look for the point of enhanced symmetry as we vary L .

To determine the location of this enhanced symmetry point, it is more convenient to work with the representation of the surface Σ where it is formed from identifications on a finite region in the hyperbolic plane. Our geometry on the pair of pants is depicted in figure 2, and is determined by the lengths L_1, L_2, L_3 of the three geodesics. This alternative representation rewrites Σ as the region in figure 2 bounded by the minimal geodesics, for some other values L'_1, L'_2, L'_3 . Unfortunately the map $L'_a(L_a)$ is not known explicitly. But in the symmetric case $L_1 = L_2 = L_3 = L$, we will have $L'_1 = L'_2 = L'_3 = L'$.

There are three parameters ψ_1, ψ_2, α_2 labeling this identification, which determine the lengths L_1, L_2, L_3 of the three geodesics. To consider the symmetric case, we need to set these three lengths to be the same; we need to determine the choice of ψ_1, ψ_2, α_2 to which this corresponds. A relatively simple way to do so is to split the fundamental region in figure 2 along the straight line $\phi = 0, \pi$. We may then think of our pair of pants as being built by gluing together two regions in the hyperbolic plane, each bounded by three geodesics; we will henceforth refer to these bounding geodesics as seams. The above split has the effect of separating every minimal closed geodesic H_a of figure 2 into a pair of segments, with each segment running between some pair of seams. This description is manifestly symmetric under permuting the boundaries when the seams are taken to be the geodesics

$$\tanh \chi \cos(\phi - \alpha'_a) = \cos \psi', \quad (\text{A.1})$$

whith $\alpha'_1 = \frac{\pi}{6}$, $\alpha'_2 = \frac{7\pi}{6}$, and $\alpha'_3 = \frac{3\pi}{2}$, see figure 8. A series of boosts relates this to the presentation in figure 2, from which we learn that

$$\cos^2 \psi = 1 - \frac{\sin^4 \psi'}{(\sin \pi/6 + \cos^2 \psi')^2} \quad (\text{A.2})$$

and thus that

$$L = 2\ell \tanh^{-1} \left[\left(1 - \frac{\sin^4 \psi'}{(\sin \pi/6 + \cos^2 \psi')^2} \right)^{1/2} \right] \quad (\text{A.3})$$

The single parameter ψ can clearly run from 0 (where $L \rightarrow \infty$) to $\pi/3$ (where $L \rightarrow 0$). For the representation of Σ as the interior region, there is a similar expression for L' as a function of the corresponding angle.

Gluing together two of these pairs of pants along the seams gives us the genus two surface corresponding to the case with all L_a equal. In the representation as a finite region, it is clear that the interesting enhanced symmetry point is when the pair of pants has a symmetry under interchanging the identified and minimal geodesics, that is when the length along the identified geodesic between two minimal geodesics in figure 8 is the same as L' . At this

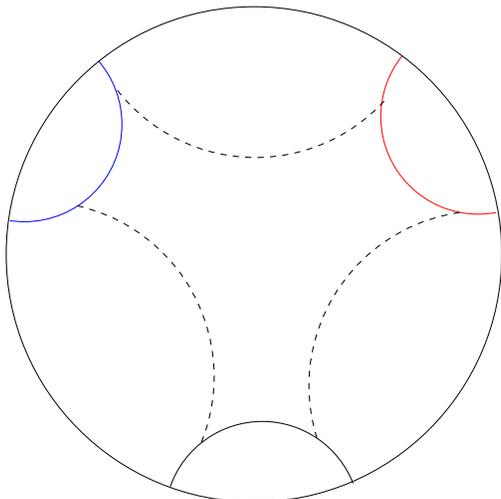


Figure 8: The symmetric representation of the pair of pants geometry; the geometry is obtained by taking two copies of the Poincaré disc and identifying corresponding geodesics in each copy. For the symmetric configuration where the horizon lengths are equal, this figure has a symmetry under rotations by $2\pi/3$. The minimal closed geodesics are made up of the dashed lines in the two copies. The symmetric point is where the length of the minimal geodesics is equal to that of the identified lines between the minimal geodesics, when the figure becomes a hyperbolic regular hexagon.

symmetric point, the genus two surface will be the Bolza surface, the most symmetric genus two surface (see e.g. [67] for further discussion of symmetric genus two surfaces).

This is most conveniently worked out in the coordinates of figure 2, where the length along the identified geodesic $\phi = \pi$ between the two minimal geodesics is simply $L_{id} = \ell_{\chi_{min}}$, where χ_{min} is the minimum value of χ along the minimal geodesic. Boosting to the coordinates of figure 8, we find this is

$$L_{id} = 2\ell \tanh^{-1} \left[\frac{\cos \pi/6}{(1 + \sin \pi/6)} \tan \psi' \right], \quad (\text{A.4})$$

so the value of ψ' at which $L_{id} = L'$ is determined by solving

$$\frac{\cos^2 \pi/6}{(1 + \sin^2 \pi/6)^2} \tan^2 \psi = 1 - \frac{\sin^4 \psi}{(\sin \pi/6 + \cos^2 \psi)^2}, \quad (\text{A.5})$$

which gives $\psi = 0.848906$, corresponding to a length $L' = 1.701$. As expected, the transition occurs at parameter values of order one. The corresponding value of L can't be explicitly determined, but it should also be order one.

B Derivation of the k -producible bound

We now derive the bound (4.29) on entropies associated with a k -producible pure n -party state. We define $\delta S(B_X) := S_{\text{sum}}(B_X) - S(B_X)$ and $X^c := [n] \setminus X$, with all other notation

as in section 4.6.

The first step will be to prove the following statement: let $X, Y \in \mathcal{C}_j^{[n]}$ label collections of systems and $K \subseteq [n]$ be a subset of size $k \leq n - 1$, then

$$\max [\delta S(B_X^K), \delta S(B_Y^K)] \geq S_{\text{sum}}(B_{X \cap Y}^K) - S_{\text{sum}}(B_{X^c \cap Y^c}^K). \quad (\text{B.1})$$

To get a feel for the inequality, consider the simplest case, when $X \cup Y = K$, and suppose for the sake of argument that $\delta S(B_X^K) = \delta S(B_Y^K) = 0$. Then $S(B_Y^K) = S(B_{Y^c}^K)$ since the state on $B_{Y \cup Y^c}^K$ is pure, so

$$S(B_Y^K) = S(B_{Y^c}^K) = S_{\text{sum}}(B_{Y^c}^K) \quad (\text{B.2})$$

$$= S_{\text{sum}}(B_X^K) - S_{\text{sum}}(B_{X \cap Y}^K) \quad (\text{B.3})$$

$$= S_{\text{sum}}(B_{X^c}^K) - S_{\text{sum}}(B_{X \cap Y}^K) \quad (\text{B.4})$$

$$= S_{\text{sum}}(B_Y^K) - 2S_{\text{sum}}(B_{X \cap Y}^K). \quad (\text{B.5})$$

Comparing the left and right hand sides of the equation, we conclude that $S_{\text{sum}}(B_{X \cap Y}^K) = 0$. Increasing $\delta S(B_X^K)$ or $\delta S(B_Y^K)$ allows $S_{\text{sum}}(B_{X \cap Y}^K)$ to also increase in a controlled way, which is the content of (B.1).

The proof of (B.1) hinges on the following easily proved but complicated-looking formula:

$$S(B_Y^K) = S_{\text{sum}}(B_Y^K) - 2S_{\text{sum}}(B_{X \cap Y}^K) + 2S_{\text{sum}}(B_{X^c \cap Y^c}^K) + \delta S(B_X^K) - \delta S(B_{X^c}^K) - \delta S(B_{Y^c}^K). \quad (\text{B.6})$$

The demonstration is just an exercise in set-theoretic manipulations and use of the fact that $S(B_Z^K) = S(B_{Z^c}^K)$.

$$S(B_Y^K) = S_{\text{sum}}(B_{Y^c}^K) - \delta S(B_{Y^c}^K) \quad (\text{B.7})$$

$$= S_{\text{sum}}(B_{Y^c}^K \cap B_X^K) + S_{\text{sum}}(B_{Y^c}^K \cap B_{X^c}^K) - \delta S(B_{Y^c}^K) \quad (\text{B.8})$$

$$= S_{\text{sum}}(B_X^K) - S_{\text{sum}}(B_{X \cap Y}^K) + S_{\text{sum}}(B_{X^c \cap Y^c}^K) - \delta S(B_{Y^c}^K) \quad (\text{B.9})$$

$$= S(B_X^K) - S_{\text{sum}}(B_{X \cap Y}^K) + S_{\text{sum}}(B_{X^c \cap Y^c}^K) - \delta S(B_X^K) - \delta S(B_{Y^c}^K) \quad (\text{B.10})$$

$$= S(B_{X^c}^K) - S_{\text{sum}}(B_{X \cap Y}^K) + S_{\text{sum}}(B_{X^c \cap Y^c}^K) - \delta S(B_X^K) - \delta S(B_{Y^c}^K) \quad (\text{B.11})$$

$$= S_{\text{sum}}(B_{X^c}^K) - S_{\text{sum}}(B_{X \cap Y}^K) + S_{\text{sum}}(B_{X^c \cap Y^c}^K) - \delta S(B_X^K) - \delta S(B_{X^c}^K) - \delta S(B_{Y^c}^K). \quad (\text{B.12})$$

But

$$S_{\text{sum}}(B_{X^c}^K) = S_{\text{sum}}(B_{X^c \cap Y}^K) + S_{\text{sum}}(B_{X^c \cap Y^c}^K) \quad (\text{B.13})$$

$$= S_{\text{sum}}(B_Y^K) - S_{\text{sum}}(B_{X \cap Y}^K) + S_{\text{sum}}(B_{X^c \cap Y^c}^K). \quad (\text{B.14})$$

(B.6) then follows by substitution into (B.12).

We will need to eliminate $\delta S(B_{X^c}^K)$ and $\delta S(B_{Y^c}^K)$ from (B.6) to get a state-independent bound. By the subadditivity of entropy, δS is nonnegative. Moreover, the monotonicity of the relative entropy function $S(\rho \parallel \sigma) = \text{tr} \rho(\log \rho - \log \sigma)$ can be used to conclude that

$$\delta S(B_Z^K) = S(\psi_Z^K \parallel \otimes_{z \in Z} \psi_z^K) \geq S(\psi_{Z'}^K \parallel \otimes_{z \in Z'} \psi_z^K) = \delta S(B_{Z'}^K), \quad (\text{B.15})$$

provided $Z' \subseteq Z$. Therefore,

$$2S_{\text{sum}}(B_{X^c \cap Y^c}^K) - \delta S(B_{X^c}^K) - \delta S(B_{Y^c}^K) \quad (\text{B.16})$$

$$\leq 2S_{\text{sum}}(B_{X^c \cap Y}^K) - 2\delta S(B_{X^c \cap Y^c}^K) \quad (\text{B.17})$$

$$= 2S(B_{X^c \cap Y^c}^K). \quad (\text{B.18})$$

$$(\text{B.19})$$

Substituting back into (B.6) leads to the inequality

$$\delta S(B_Y^K) \geq 2S_{\text{sum}}(B_{X \cap Y}) - 2S_{\text{sum}}(B_{X^c \cap Y^c}^K) + \delta S(B_X^K), \quad (\text{B.20})$$

which means that

$$\delta S(B_X^K) + \delta S(B_Y^K) \geq \max[\delta S(B_X^K), \delta S(B_Y^K)] \quad (\text{B.21})$$

$$\geq \max[\delta S(B_X^K), 2S_{\text{sum}}(B_{X \cap Y}^K) - 2S_{\text{sum}}(B_{X^c \cap Y^c}^K) - \delta S(B_X^K)] \quad (\text{B.22})$$

$$\geq S_{\text{sum}}(B_{X \cap Y}^K) - S_{\text{sum}}(B_{X^c \cap Y^c}^K), \quad (\text{B.23})$$

which subsumes inequality (B.1). The third line is just a consequence of optimizing the unknown nonnegative quantity $\delta S(B_X^K)$.

Up to now, we have been studying individual states that are pure on B^K . To get a bound applicable to all k -producible pure states, it will be necessary to remove the dependence on K , which we will do by summing (B.23) over X, Y and K . On the left hand side, using that B_X is isometrically related to $\otimes_{K \in \mathcal{C}_k^{[n]}} \otimes_{x \in X} B_x^K$, we can calculate

$$\sum_{X, Y \in \mathcal{C}_j^{[n]}} \sum_{K \in \mathcal{C}_k^{[n]}} [\delta S(B_X^K) + \delta S(B_Y^K)] = \sum_{X, Y \in \mathcal{C}_j^{[n]}} [\delta S(B_X) + \delta S(B_Y)] \quad (\text{B.24})$$

$$= 2 \binom{n}{j} \sum_{X \in \mathcal{C}_j^{[n]}} \delta S(B_X). \quad (\text{B.25})$$

The evaluation of the right hand side of (B.23) is significantly more involved. As a first observation, recall that for generic states, $\delta S(B_X)$ is zero only if $|X| \leq n/2$ so we are only concerned with $j = |X| = |Y| \leq n/2$, in which case $|X \cap Y| \leq |X^c \cap Y^c|$. Summing (B.23) over all K , however, will typically give a trivial bound under those circumstances because of negative contributions. Instead, since $\delta S(B_X^K) + \delta S(B_Y^K)$ is nonnegative, (B.25) can more fruitfully be bounded below by

$$\sum_{X, Y \in \mathcal{C}_j^{[n]}} \sum_{K \in \mathcal{C}_k^{[n]}} I[(X, Y, K) \in \mathcal{G}] (S_{\text{sum}}(B_{X \cap Y}^K) - S_{\text{sum}}(B_{X^c \cap Y^c}^K)), \quad (\text{B.26})$$

where $I[\cdot]$ denotes the indicator function ($I[\cdot] = 1$ when the condition is satisfied, and is zero otherwise) and we have the freedom to choose \mathcal{G} as we please in order to get the best bound

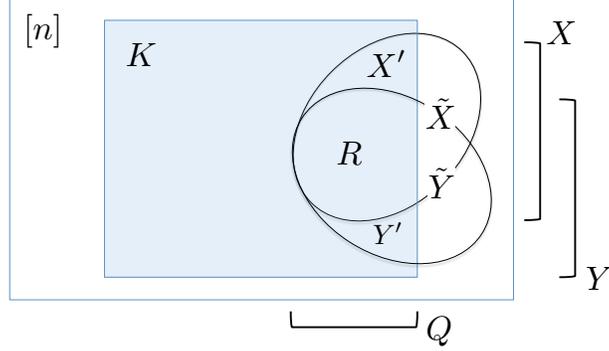


Figure 9: Decomposition of the sets X and Y used in the evaluation of (B.26).

possible. Let $Q = (X \cup Y) \cap K$ and $R = X \cap Y \cap K$. We will take \mathcal{G} to be the set such that $|Q| \in \mathcal{Q}$ and $|R| \in \mathcal{R}$ for choices of \mathcal{Q} and \mathcal{R} to be determined. To facilitate the calculation, break X and Y each into three parts, namely their intersection R , their portions in K but not in R , and the rest, as depicted in figure 9. In terms of those definitions, (B.26) can be rewritten as

$$\begin{aligned} & \sum_{(X,Y,K) \in \mathcal{G}} [S_{\text{sum}}(B_{X \cap Y}^K) - S_{\text{sum}}(B_{X^c \cap Y^c}^K)] \quad (\text{B.27}) \\ = & \sum_{K \in \mathcal{C}_k^{[n]}} \sum_{Q \in \mathcal{C}_q^K} \sum_{R \in \mathcal{C}_r^Q} \sum_{X' \in \mathcal{D}} \sum_{\tilde{X} \in \mathcal{C}_{j-r-|X'|}^{[n] \setminus K}} \sum_{\tilde{Y} \in \mathcal{C}_{j-r-q+|X'|}^{[n] \setminus K}} [S_{\text{sum}}(B_{X \cap Y}^K) - S_{\text{sum}}(B_{X^c \cap Y^c}^K)] \quad (\text{B.28}) \end{aligned}$$

where X and Y are the disjoint unions $R \cup X' \cup \tilde{X}$ and $R \cup Y' \cup \tilde{Y}$, respectively, while $\mathcal{D} = \{X' \subseteq Q \setminus R : \max(0, q-j) \leq |X'| \leq \min(j-r, q-r)\}$. The constraints on the size of X' arise from the requirement that $|X'| + |Y'| + r = q$ but that $|X'|, |Y'| \leq j-r$.

Since the summand is independent of X', \tilde{X} and \tilde{Y} , the three rightmost sums can be evaluated, yielding a multiplicative factor of

$$C_{jknqr} = \sum_{l=\max(0, q-j)}^{\min(q-r, j-r)} \binom{q-r}{l} \binom{n-k}{j-r-l} \binom{n-k}{j-r-q+l}, \quad (\text{B.29})$$

with l playing the role of $|X'|$. Next, we compute the sum of $S_{\text{sum}}(B_{X \cap Y}^K)$:

$$\sum_{K \in \mathcal{C}_k^{[n]}} \sum_{Q \in \mathcal{C}_q^K} \sum_{R \in \mathcal{C}_r^Q} S_{\text{sum}}(B_{X \cap Y}^K) = \sum_{K \in \mathcal{C}_k^{[n]}} \sum_{R \in \mathcal{C}_r^K} \sum_{Q' \in \mathcal{C}_{q-r}^{K \setminus R}} S_{\text{sum}}(B_R^K) \quad (\text{B.30})$$

$$= \binom{k-r}{q-r} \sum_{K \in \mathcal{C}_k^{[n]}} \sum_{R \in \mathcal{C}_r^K} S_{\text{sum}}(B_R^K) \quad (\text{B.31})$$

$$= \binom{k-r}{q-r} \sum_{x \in [n]} \sum_{K' \in \mathcal{C}_{k-1}^{[n] \setminus \{x\}}} \sum_{R' \in \mathcal{C}_{r-1}^{K'}} S(B_x^K) \quad (\text{B.32})$$

$$= \binom{k-r}{q-r} \binom{k-1}{r-1} \sum_{x \in [n]} S(B_x) \quad (\text{B.33})$$

$$= \binom{k-1}{k-q, q-r, r-1} \sum_{x \in [n]} S(B_x), \quad (\text{B.34})$$

In simplifying the sums, we have used the fact that $S(B_x) = \sum_K S(B_x^K)$. The contribution to (B.28) from the $S_{\text{sum}}(B_{X^c \cap Y^c}^K)$ terms can then be disposed of with a similar calculation.

$$\sum_{K \in \mathcal{C}_k^{[n]}} \sum_{Q \in \mathcal{C}_q^K} \sum_{R \in \mathcal{C}_r^Q} S_{\text{sum}}(B_{Q^c}^K) = \binom{q}{r} \sum_{K \in \mathcal{C}_k^{[n]}} \sum_{Q^c \in \mathcal{C}_{k-q}^{K^c}} S_{\text{sum}}(B_{Q^c}^K) \quad (\text{B.35})$$

$$= \binom{q}{r} \sum_{x \in [n]} \sum_{K' \in \mathcal{C}_{k-1}^{[n] \setminus \{x\}}} \sum_{\tilde{Q}^c \in \mathcal{C}_{k-q-1}^{K'^c}} S(B_x^{\{x\} \cup K'}) \quad (\text{B.36})$$

$$= \binom{q}{r} \binom{k-1}{k-q-1} \sum_{x \in [n]} S(B_x) \quad (\text{B.37})$$

$$= \binom{k-1}{k-q-1, q-r, r} \sum_{x \in [n]} S(B_x). \quad (\text{B.38})$$

Substituting leads to the conclusion that (B.28) is equal to

$$C_{jknqr} \binom{k-1}{k-q, q-r, r-1} \left(1 - \frac{k-q}{r}\right) \sum_{x \in [n]} S(B_x). \quad (\text{B.39})$$

While that completes the proof of the bound, the conclusion can be expressed slightly more conveniently in terms of

$$\tilde{C}_{jknqr} = \sum_{l=\max(0, q-j)}^{\min(q-r, j-r)} \binom{k-1}{k-q, q-r-l, l, r-1} \binom{n-k}{j-r-l} \binom{n-k}{j-r-q+l}. \quad (\text{B.40})$$

We have shown that

$$2 \binom{n}{j} \sum_{X \in \mathcal{C}_j^{[n]}} \delta S(B_X) \geq \sum_{q \in \mathcal{Q}} \sum_{r \in \mathcal{R}} \tilde{C}_{jknqr} \left(1 - \frac{k-q}{r}\right) \sum_{x \in [n]} S(B_x). \quad (\text{B.41})$$

Note that the formula only makes sense when q and r are chosen to be consistent with j, n and k , which requires that $r \leq q \leq k$ and $q \leq 2j - r$. This is a general inequality that must be satisfied for every k -producible state on n factors. The strongest bound occurs when \mathcal{Q}, \mathcal{R} are selected to contain precisely those q, r making positive contributions, which is the choice made to produce the plots in the main text.

References

- [1] J. Maldacena, *Eternal black holes in anti-de Sitter*, *Journal of High Energy Physics* **2003** (2003), no. 04 021, [[hep-th/0106112](#)].
- [2] M. Van Raamsdonk, *Building up spacetime with quantum entanglement*, *General Relativity and Gravitation* **42** (2010), no. 10 2323–2329, [[arXiv:1005.3035](#)].
- [3] B. Czech, J. L. Karczmarek, F. Nogueira, and M. Van Raamsdonk, *Rindler quantum gravity*, *Classical and Quantum Gravity* **29** (2012), no. 23 235025, [[arXiv:1206.1323](#)].
- [4] J. Maldacena and L. Susskind, *Cool horizons for entangled black holes*, *Fortschritte der Physik* **61** (2013), no. 9 781–811, [[arXiv:1306.0533](#)].
- [5] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Physical review letters* **96** (2006), no. 18 181602, [[hep-th/0603001](#)].
- [6] V. E. Hubeny, M. Rangamani, and T. Takayanagi, *A covariant holographic entanglement entropy proposal*, *Journal of High Energy Physics* **2007** (2007), no. 07 062, [[hep-th/0705.0016](#)].
- [7] A. Lewkowycz and J. Maldacena, *Generalized gravitational entropy*, *Journal of High Energy Physics* **2013** (2013), no. 8 1–29, [[arXiv:1304.4926](#)].
- [8] E. Bianchi and R. C. Myers, *On the Architecture of Spacetime Geometry*, [arXiv:1212.5183](#).
- [9] V. Balasubramanian, B. Czech, B. D. Chowdhury, and J. de Boer, *The entropy of a hole in spacetime*, *JHEP* **1310** (2013) 220, [[arXiv:1305.0856](#)].
- [10] V. Balasubramanian, B. D. Chowdhury, B. Czech, J. de Boer, and M. P. Heller, *A hole-ographic spacetime*, [arXiv:1310.4204](#).
- [11] R. C. Myers, J. Rao, and S. Sugishita, *Holographic holes in higher dimensions*, [arXiv:1403.3416](#).
- [12] B. Swingle, *Entanglement renormalization and holography*, *Physical Review D* **86** (2012), no. 6 065007, [[arXiv:0905.1317](#)].

- [13] M. Nozaki, T. Numasawa, A. Prudenziati, and T. Takayanagi, *Dynamics of Entanglement Entropy from Einstein Equation*, *Phys.Rev.* **D88** (2013), no. 2 026012, [[arXiv:1304.7100](#)].
- [14] N. Lashkari, M. B. McDermott, and M. Van Raamsdonk, *Gravitational Dynamics From Entanglement "Thermodynamics"*, *JHEP* **1404** (2014) 195, [[arXiv:1308.3716](#)].
- [15] T. Faulkner, M. Guica, T. Hartman, R. C. Myers, and M. Van Raamsdonk, *Gravitation from Entanglement in Holographic CFTs*, *JHEP* **1403** (2014) 051, [[arXiv:1312.7856](#)].
- [16] B. Swingle and M. Van Raamsdonk, *Universality of Gravity from Entanglement*, [arXiv:1405.2933](#).
- [17] B. S. Cirel'son, *Quantum generalizations of Bell's inequality*, *Letters in Mathematical Physics* **4** (1980), no. 2 93–100.
- [18] D. M. Greenberger, M. A. Horne, A. Shimony, and A. Zeilinger, *Bells theorem without inequalities*, *Am. J. Phys* **58** (1990), no. 12 1131–1143.
- [19] K. Krasnov, *Holography and Riemann surfaces*, *Adv.Theor.Math.Phys.* **4** (2000) 929–979, [[hep-th/0005106](#)].
- [20] K. Krasnov, *Black hole thermodynamics and Riemann surfaces*, *Class.Quant.Grav.* **20** (2003) 2235–2250, [[gr-qc/0302073](#)].
- [21] K. Skenderis and B. C. van Rees, *Holography and wormholes in 2+ 1 dimensions*, *Communications in Mathematical Physics* **301** (2011), no. 3 583–626, [[arXiv:0912.2090](#)].
- [22] L. Susskind, *Addendum to Computational Complexity and Black Hole Horizons*, [arXiv:1403.5695](#).
- [23] P. Hayden, M. Headrick, and A. Maloney, *Holographic mutual information is monogamous*, *Physical Review D* **87** (2013), no. 4 046003, [[arXiv:1107.2940](#)].
- [24] H. Gharibyan and R. F. Penna, *Are entangled particles connected by wormholes? Support for the ER=EPR conjecture from entropy inequalities*, [arXiv:1308.0289](#).
- [25] D. R. Brill, *Multi-black-hole geometries in (2+ 1)-dimensional gravity*, *Physical Review D* **53** (1996), no. 8 R4133, [[gr-qc/9511022](#)].
- [26] S. Aminneborg, I. Bengtsson, D. Brill, S. Holst, and P. Peldan, *Black holes and wormholes in (2+1)-dimensions*, *Class. Quant. Grav.* **15** (1998) 627–644, [[gr-qc/9707036](#)].
- [27] D. Brill, *Black holes and wormholes in 2+ 1 dimensions*, in *Mathematical and quantum aspects of relativity and cosmology*, pp. 143–179. Springer, 2000. [gr-qc/9904083](#).

- [28] S. Aminneborg, I. Bengtsson, and S. Holst, *A spinning anti-de Sitter wormhole*, *Class. Quant. Grav.* **16** (1999) 363–382, [[gr-qc/9805028](#)].
- [29] V. Vedral, M. B. Plenio, M. A. Rippin, and P. L. Knight, *Quantifying entanglement*, *Phys. Rev. Lett.* **78** (Mar, 1997) 2275–2279, [[quant-ph/9702027](#)].
- [30] C. H. Bennett, S. Popescu, D. Rohrlich, J. A. Smolin, and A. V. Thapliyal, *Exact and asymptotic measures of multipartite pure-state entanglement*, *Physical Review A* **63** (2000), no. 1 012307, [[quant-ph/9908073](#)].
- [31] H. Barnum and N. Linden, *Monotones and invariants for multi-particle quantum states*, *Journal of Physics A: Mathematical and General* **34** (2001), no. 35 6787, [[quant-ph/0103155](#)].
- [32] J. Eisert and D. Gross, *Multi-particle entanglement*, [[quant-ph/0505149](#)].
- [33] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, *Quantum entanglement*, *Reviews of Modern Physics* **81** (2009), no. 2 865, [[quant-ph/0702225](#)].
- [34] C. H. Bennett, H. J. Bernstein, S. Popescu, and B. Schumacher, *Concentrating partial entanglement by local operations*, *Physical Review A* **53** (1996), no. 4 2046, [[quant-ph/9511030](#)].
- [35] W. Dür, G. Vidal, and J. I. Cirac, *Three qubits can be entangled in two inequivalent ways*, *Physical Review A* **62** (2000), no. 6 062314, [[quant-ph/0005115](#)].
- [36] D. N. Page, *Average entropy of a subsystem*, *Physical Review Letters* **71** (1993), no. 9 1291.
- [37] M. Banados, C. Teitelboim, and J. Zanelli, *Black hole in three-dimensional spacetime*, *Physical Review Letters* **69** (1992), no. 13 1849, [[hep-th/9204099](#)].
- [38] D. Marolf and A. C. Wall, *Eternal Black Holes and Superselection in AdS/CFT*, *Class.Quant.Grav.* **30** (2013) 025001, [[arXiv:1210.3590](#)].
- [39] S. H. Shenker and D. Stanford, *Black holes and the butterfly effect*, *JHEP* **1403** (2014) 067, [[arXiv:1306.0622](#)].
- [40] S. H. Shenker and D. Stanford, *Multiple Shocks*, [[arXiv:1312.3296](#)].
- [41] J. Louko and D. Marolf, *Inextendible Schwarzschild black hole with a single exterior: How thermal is the Hawking radiation?*, *Phys.Rev.* **D58** (1998) 024007, [[gr-qc/9802068](#)].
- [42] M. Headrick, V. Hubeny, M. Rangamani, and A. Lawrence, “Causality and holographic entanglement entropy.” Work in progress described at New Frontiers in Dynamical Gravity, Cambridge, 2014.

- [43] D. Marolf and M. Rangamani, *Holography and the topology of causal shadows*, . to appear.
- [44] B. Maskit, *Kleinian Groups*. Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen mit besonderer Berücksichtigung der Anwendungsgebiete. Springer, 1987.
- [45] J. M. Maldacena and A. Strominger, *AdS(3) black holes and a stringy exclusion principle*, *JHEP* **9812** (1998) 005, [[hep-th/9804085](#)].
- [46] X. Yin, *On Non-handlebody Instantons in 3D Gravity*, *JHEP* **0809** (2008) 120, [[arXiv:0711.2803](#)].
- [47] M. Seppala and T. Sorvali, *Geometry of Riemann surfaces and Teichmuller spaces*. Elsevier, 2011.
- [48] T. Faulkner, *The Entanglement Renyi Entropies of Disjoint Intervals in AdS/CFT*, [arXiv:1303.7221](#).
- [49] T. Hartman, *Entanglement Entropy at Large Central Charge*, [arXiv:1303.6955](#).
- [50] R. F. Werner, *Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model*, *Physical Review A* **40** (1989), no. 8 4277–4281.
- [51] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, *Mixed-state entanglement and quantum error correction*, *Physical Review A* **54** (1996), no. 5 3824, [[quant-ph/9604024](#)].
- [52] A. C. Wall, *Maximin surfaces, and the strong subadditivity of the covariant holographic entanglement entropy*, [arXiv:1211.3494](#).
- [53] I. Devetak and A. Winter, *Distillation of secret key and entanglement from quantum states*, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science* **461** (2005), no. 2053 207–235, [[quant-ph/0306078](#)].
- [54] M. Horodecki, J. Oppenheim, and A. Winter, *Quantum state merging and negative information*, *Communications in Mathematical Physics* **269** (2007), no. 1 107–136, [[quant-ph/0512247](#)].
- [55] M. Horodecki, P. Horodecki, and R. Horodecki, *Limits for entanglement measures*, *Physical Review Letters* **84** (2000), no. 9 2014, [[quant-ph/9908065](#)].
- [56] P. M. Hayden, M. Horodecki, and B. M. Terhal, *The asymptotic entanglement cost of preparing a quantum state*, *Journal of Physics A: Mathematical and General* **34** (2001), no. 35 6891, [[quant-ph/0008134](#)].

- [57] P. Hayden, D. W. Leung, and A. Winter, *Aspects of generic entanglement*, *Communications in Mathematical Physics* **265** (2006), no. 1 95–117, [[quant-ph/0407049](#)].
- [58] M. Christandl and A. Winter, *Squashed entanglement: An additive entanglement measure*, *Journal of mathematical physics* **45** (2004), no. 3 829–840, [[quant-ph/0308088](#)].
- [59] O. Gühne, G. Tóth, and H. J. Briegel, *Multipartite entanglement in spin chains*, *New Journal of Physics* **7** (2005), no. 1 229, [[quant-ph/0502160](#)].
- [60] D. Marolf and J. Polchinski, *Gauge/Gravity Duality and the Black Hole Interior*, *Phys.Rev.Lett.* **111** (2013) 171301, [[arXiv:1307.4706](#)].
- [61] V. Balasubramanian, M. Berkooz, S. F. Ross, and J. Simon, *Black Holes, Entanglement and Random Matrices*, [arXiv:1404.6198](#).
- [62] S. T. Belinschi, B. Collins, and I. Nechita, “Almost one bit violation for the additivity of the minimum output entropy.” 2013.
- [63] P. Hayden and N. Thomas, “Random tensor networks and multiboundary black holes.” To appear., 2014.
- [64] A. R. Steif, *Time symmetric initial data for multibody solutions in three-dimensions*, *Phys.Rev.* **D53** (1996) 5527–5532, [[gr-qc/9511053](#)].
- [65] A. Maloney and E. Witten, *Quantum Gravity Partition Functions in Three Dimensions*, *JHEP* **1002** (2010) 029, [[arXiv:0712.0155](#)].
- [66] R. Dijkgraaf, J. M. Maldacena, G. W. Moore, and E. P. Verlinde, *A Black hole Farey tail*, [hep-th/0005003](#).
- [67] M. G. Katz, *Systolic Geometry and Topology*. AMS, 2007.