

 Open access • Journal Article • DOI:10.1007/JHEP08(2013)092

Holographic entanglement plateaux — [Source link](#)

[Veronika E. Hubeny](#), [Henry Maxfield](#), [Mukund Rangamani](#), [Erik Tonni](#)

Institutions: [International School for Advanced Studies](#)

Published on: 17 Jun 2013 - [Journal of High Energy Physics](#) (Springer Berlin Heidelberg)

Topics: [Squashed entanglement](#), [Quantum entanglement](#), [Entropy \(arrow of time\)](#) and [AdS/CFT correspondence](#)

Related papers:

- [Holographic Derivation of Entanglement Entropy from the anti de Sitter Space/Conformal Field Theory Correspondence](#)
- [A covariant holographic entanglement entropy proposal](#)
- [Aspects of Holographic Entanglement Entropy](#)
- [Generalized gravitational entropy](#)
- [Holographic proof of the strong subadditivity of entanglement entropy](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/holographic-entanglement-plateaux-4oo1q2hn2t>

Durham Research Online

Deposited in DRO:

11 December 2013

Version of attached file:

Other

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Hubeny, Veronika E. and Maxfield, Henry and Rangamani, Mukund and Tonni, Erik (2013) 'Holographic entanglement plateaux.', *Journal of high energy physics.*, 2013 (8). p. 92.

Further information on publisher's website:

[http://dx.doi.org/10.1007/JHEP08\(2013\)092](http://dx.doi.org/10.1007/JHEP08(2013)092)

Publisher's copyright statement:

© SISSA 2013. Published by Springer on behalf of International School for Advanced Studies (SISSA - Trieste, Italy). The final publication is available at Springer via [http://dx.doi.org/10.1007/JHEP08\(2013\)092](http://dx.doi.org/10.1007/JHEP08(2013)092).

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

PREPARED FOR SUBMISSION TO JHEP

Holographic entanglement plateaux

Veronika E. Hubeny^a, Henry Maxfield^a, Mukund Rangamani^a & Erik Tonni^b

^a*Centre for Particle Theory & Department of Mathematical Sciences,
Science Laboratories, South Road, Durham DH1 3LE, UK.*

^b*SISSA and INFN, via Bonomea 265, 34136, Trieste, Italy*

E-mail: veronika.hubeny@durham.ac.uk, h.d.maxfield@durham.ac.uk,
mukund.rangamani@durham.ac.uk, erik.tonni@sissa.it

ABSTRACT: We consider the entanglement entropy for holographic field theories in finite volume. We show that the Araki-Lieb inequality is saturated for large enough subregions, implying that the thermal entropy can be recovered from the knowledge of the region and its complement. We observe that this actually is forced upon us in holographic settings due to non-trivial features of the causal wedges associated with a given boundary region. In the process, we present an infinite set of extremal surfaces in Schwarzschild-AdS geometry anchored on a given entangling surface. We also offer some speculations regarding the homology constraint required for computing holographic entanglement entropy.

KEYWORDS: AdS-CFT correspondence, Entanglement entropy

Contents

1	Introduction	1
2	Generic behaviour of holographic $S_{\mathcal{A}}(\alpha)$	5
3	Entanglement entropy in 1 + 1 dimensional CFTs	6
4	Entanglement entropy in higher dimensional field theories	11
4.1	Entanglement plateaux in $d > 2$	12
4.2	Sub-dominant saddles: folds in minimal surfaces	15
5	Causality and holographic entanglement	17
6	Minimal versus extremal surfaces and the homology constraint	20
7	Discussion	25
A	Mean Curvature Flow	28
A.1	Some geometry	28
A.2	The algorithm	30
B	Near-horizon behaviour of minimal surfaces	32
B.1	The linear regime	32
B.2	Near-pole flat space regime	33
B.3	Matching asymptotics	33
B.4	Validity	34
B.5	Relation to θ_{∞}	35

1 Introduction

Entanglement is one of the most non-classical features of quantum mechanics and much effort has been expended in trying to figure out a quantitative measure for the amount of entanglement

inherent in a quantum state. In local quantum field theories one such measure is provided by the entanglement entropy $S_{\mathcal{A}}$ associated with a specified (spacelike) region \mathcal{A} located on a Cauchy slice Σ . This quantity is defined to be the von Neumann entropy of the reduced density matrix $\rho_{\mathcal{A}}$ obtained by integrating out the degrees of freedom in the complementary region \mathcal{A}^c . Famously the entanglement entropy is UV divergent, with the leading divergence being given by the area of the entangling surface $\partial\mathcal{A}$ (the boundary of \mathcal{A}). However, the finite part of the entanglement entropy contains non-trivial information about the quantum state and in certain cases serves as a novel order parameter. As usual, it is this finite part which we shall consider.

Our interest is in understanding the behaviour of entanglement entropy for field theories defined on compact spatial geometries; so we take Σ to be a compact Riemannian manifold (typically \mathbf{S}^{d-1}). One of the questions which we wish to address is whether the entanglement entropy for a fixed state of the field on Σ is a smooth function of the (geometrical attributes of the) region \mathcal{A} . We envisage considering a family of *smooth* regions \mathcal{A}_α characterized by some parameter α , which is a proxy for the relative size of $\mathcal{A} \subset \Sigma$. The main question we want to ask is whether $S_{\mathcal{A}}(\alpha)$ is smooth under changing α .

For finite systems, analogy with statistical mechanics suggests that this would indeed be true. There should be no room for any discontinuity in $S_{\mathcal{A}}(\alpha)$ since the reduced density matrix will change analytically with α . The place where this is expected to break down is when we attain some analog of the thermodynamic limit, i.e., when the number of degrees of freedom involved gets large. So a natural place to look for non-smooth behaviour is in the dynamics of large N field theories on compact spacetimes in the planar limit. This naturally motivates the study of field theories which can be captured holographically by gravitational dynamics using the gauge/gravity correspondence. We will refer to these holographic field theories as large c (central charge) theories; for conventional non-abelian gauge theories $c \sim N^2$.

In past few years much progress has been made in understanding entanglement entropy in large N field theories (at strong coupling) thanks to the seminal work of Ryu & Takayanagi [RT] [1, 2], who gave a very simple geometric prescription for computing $S_{\mathcal{A}}$ for static states in terms of the area of a bulk minimal surface anchored on $\partial\mathcal{A}$. This prescription was extended by Hubeny, Rangamani, and Takayanagi [HRT] to arbitrary time-dependent states in [3] where one considers extremal surfaces, which can also be related to light-sheets discussed in covariant entropy bound context [4]. There have been many attempts to derive the prescription from first principles; the first was made in [5] which was critically examined in [6]. More recently, [7] gave a nice argument deriving the prescription for a special class of states (conformally invariant vacuum) and spherically symmetric regions, by converting the reduced density matrix $\rho_{\mathcal{A}}$ to a thermal density matrix. A local version of this argument has been made recently in [8] and together with the results of [9, 10] goes quite a way in establishing the holographic prescription of [1] for static states. Given these developments, it is apposite to take stock of the implications of the holographic entanglement entropy prescription for the question we have in mind.

Let us first record some basic facts about $S_{\mathcal{A}}$ which we will use to investigate the behaviour of

$S_{\mathcal{A}}(\alpha)$. It is a well known fact that when the total state of the field theory is pure, the entanglement entropy of a given region and its complement are the same: $S_{\mathcal{A}} = S_{\mathcal{A}^c}$. This however ceases to be true when the entire system is itself in a density matrix ρ_{Σ} . To measure the deviation from purity of the system we could monitor the difference $\delta S_{\mathcal{A}} = S_{\mathcal{A}} - S_{\mathcal{A}^c}$. From our perspective this quantity has some useful advantages. Firstly, it is finite since the divergent contributions which are given in terms of intrinsic and extrinsic geometry of the entangling surface $\partial\mathcal{A}$ cancel. Secondly, it is bounded from above by the von Neumann entropy of the entire density matrix ρ_{Σ} , by the Araki-Lieb inequality [11]:

$$|\delta S_{\mathcal{A}}| \equiv |S_{\mathcal{A}} - S_{\mathcal{A}^c}| \leq S_{\mathcal{A} \cup \mathcal{A}^c} = S_{\rho_{\Sigma}}. \quad (1.1)$$

So one way to phrase our original question is to ask whether $\delta S_{\mathcal{A}}(\alpha)$ is a smooth function of $\alpha \in [0, 1]$ which we take to be a suitable function of the ratio of $\text{Vol}(\mathcal{A})/\text{Vol}(\Sigma)$ such that $\delta S_{\mathcal{A}}$ is an odd-function around $\alpha = \frac{1}{2}$ and becomes the total entropy $\pm S_{\rho_{\Sigma}}$ at $\alpha = 0, 1$. The issue we want to focus on is whether there is any discontinuity either in $S_{\mathcal{A}}(\alpha)$ or its derivatives as we vary α . We will argue here that for large c field theories $S_{\mathcal{A}}(\alpha)$ (given by the RT prescription) has to be continuous for time-independent (static) density matrices ρ_{Σ} . However, there can be non-trivial behaviour in $\partial_{\alpha} S_{\mathcal{A}}(\alpha)$: we will exhibit explicit examples where the function $S_{\mathcal{A}}(\alpha)$ is continuous but fails to be differentiable.^{1,2} It is important to distinguish this from situations where the total density matrix itself varies; one can certainly have entanglement entropy of a fixed-size region which is discontinuous e.g. as a function of temperature [13, 14], as is quite familiar from the simpler example of Hawking-Page transition. Here we fix the total density matrix (which has the bulk equivalent of fixing the spacetime) and consider the entanglement entropy as function of the region.

To understand the potential issues, we need to explain one key feature of the holographic prescription of [1, 3]. To compute $S_{\mathcal{A}}$, we find extremal surfaces³ $\mathfrak{E}_{\mathcal{A}}$ in the bulk spacetime \mathcal{M} which are anchored on $\partial\mathcal{A}$; in the asymptotically AdS spacetimes we consider, we demand that $\partial\mathfrak{E}_{\mathcal{A}} = \partial\mathcal{A}$. However, there can be multiple such surfaces in a given spacetime. We are instructed to restrict attention to extremal surfaces $\mathfrak{E}_{\mathcal{A}}$ which are homologous to the boundary region \mathcal{A} under consideration [18] and from the set of such surfaces pick the one with smallest area. To wit,

$$S_{\mathcal{A}} = \min_X \frac{\text{Area}(\mathfrak{E}_{\mathcal{A}})}{4G_N}, \quad X = \mathfrak{E} : \begin{cases} \partial\mathfrak{E}_{\mathcal{A}} \equiv \mathfrak{E}_{\mathcal{A}}|_{\partial\mathcal{M}} = \partial\mathcal{A} \\ \exists \mathcal{R} \subset \mathcal{M} : \partial\mathcal{R} = \mathfrak{E}_{\mathcal{A}} \cup \mathcal{A} \end{cases} \quad (1.2)$$

¹ In fact, as we indicate in §7, in time dependent examples the situation may be much more intricate and will be discussed elsewhere [12].

² Earlier studies of holographic entanglement entropy have indeed revealed examples where $S(\alpha)$ undergoes a first order phase transition, i.e., $\partial_{\alpha} S(\alpha)$ is discontinuous due to change in the nature of the minimal surfaces [15, 16]. For example the recent analysis of [17] exhibits this phenomenon arising due to the minimal surface changing topology (even in causally trivial spacetimes) as we vary α .

³ We use the notation $\mathfrak{E}_{\mathcal{A}}$ for generic extremal surfaces and indicate minimal surfaces relevant for static geometries by $\mathfrak{M}_{\mathcal{A}}$. While we review the various assertions in terms of extremal surfaces we will for the most part (until §6) only consider minimal surfaces in this paper.

where the region \mathcal{R} is a bulk co-dimension one smooth surface (in the $d+1$ dimensional spacetime \mathcal{M}) which is bounded by the extremal surface $\mathfrak{E}_{\mathcal{A}}$ and the region \mathcal{A} on the boundary.

It was appreciated already in [18] that the homology constraint is crucial for the Araki-Lieb inequality to be satisfied. The issue was further elaborated in [19] where 1+1 dimensional CFTs on a torus were considered (see also [20] for a recent discussion). In particular, its effects are most acutely felt when ρ_{Σ} is a density matrix, for then we anticipate the bulk spacetime to have a horizon [21].⁴ In static spacetimes this implies a non-trivial topology in the bulk when restricted to a constant time slice, which can be easily intuited by considering the Euclidean section of the geometry.

The homology constraint, being non-local from the boundary perspective, allows non-trivial behaviour in the nature of the extremal surfaces that are admissible for the problem. Indeed as we will show explicitly, there are many simple examples where one has multiple extremal surfaces and only some of them are homologous to the boundary region in question. In fact, the most bizarre aspect of our analysis is that for certain choices of boundary regions there are no connected minimal surfaces satisfying the homology constraint: one is forced into considering disconnected surfaces.⁵ Multiply connected extremal surfaces in turn imply that one can have novel behaviour in $S_{\mathcal{A}}$ or equivalently in $\delta S_{\mathcal{A}}$; we go on to show that in static spacetimes these indicate that $\delta S_{\mathcal{A}}(\alpha)$ is continuous but not differentiable, exhibiting explicit examples involving global AdS black hole geometries. We argue that the lack of differentiability is the worst it gets for $\delta S_{\mathcal{A}}$ in static spacetimes, proving that $\delta S_{\mathcal{A}}$ has to be a continuous function of α . Furthermore, when we are forced onto the branch of disconnected extremal surfaces, it is easy to establish that $\delta S_{\mathcal{A}} = S_{\rho_{\Sigma}}$, i.e., the Araki-Lieb inequality is saturated; we refer to this phenomenon as *entanglement plateau*.⁶

While in simple examples one can establish entanglement plateaux by explicit construction, it is interesting to examine when it should happen on general grounds. Curiously, it is easy to provide a bound, though in a somewhat roundabout manner using machinery outside the extremal surface technology. The necessary concept is geometric and has to do with the behaviour of the causal wedge associated with \mathcal{A} . These objects were studied in the context of ‘causal holographic information’ in [22] from a very different perspective (the motivation being to characterize the minimal amount of holographic information in the reduced density matrix). Using the topology of the causal wedge one can establish criteria for when the extremal surfaces $\mathfrak{E}_{\mathcal{A}}$ become disconnected. The precise statement and its proof will appear elsewhere [23], but we will flesh out the physical aspects of the argument in what follows. Suffice to say for now that we find the interplay between causality and extremal surfaces extremely intriguing and believe that it points to some yet to be fathomed facet of holography.

⁴ This statement presumes that the von Neumann entropy of ρ_{Σ} scales like the central charge of the field theory.

⁵ The exchange of dominance between connected and disconnected surfaces in confining backgrounds for field theories on non-compact geometries $\mathbb{R}^{1,d-2} \times \mathbf{S}^1$ has been well studied in the context of holographic entanglement entropy, cf., [15, 16] for initial work. We will however be considering field theories on compact spatial volumes.

⁶ This phenomenon like many others encountered in holographic duals is a feature of large c theories. At finite central charge we cannot have any sharp plateaux; we thank Hong Liu for discussions on this issue.

The outline of the paper is as follows: we begin in §2 with a discussion of the general behaviour of the entanglement entropy $S_{\mathcal{A}}$ as function of the size parameter α and argue that in the holographic context, $S_{\mathcal{A}}(\alpha)$ and consequentially $\delta S_{\mathcal{A}}(\alpha)$ should be continuous functions. We then proceed to see the explicit behaviour of entanglement entropy in 1+1 dimensional CFTs on a torus in §3 and in higher dimensional CFTs in §4. After displaying explicit examples of the entanglement plateaux, we outline the connection with causal wedges in §5. We then step back to compare the prescriptions involving minimal (RT) versus extremal (HRT) surfaces and relatedly the role of the homology requirement in §6 and conclude with a discussion of open issues in §7. Some technical aspects of finding the minimal surfaces are relegated to the Appendices.

2 Generic behaviour of holographic $S_{\mathcal{A}}(\alpha)$

We begin our discussion by explaining the continuity of $S_{\mathcal{A}}(\alpha)$. Importantly, we focus on the RT prescription for static configurations, where the entire problem can be formulated on a d dimensional Riemannian bulk geometry with $(d - 1)$ -dimensional boundary (in the conformal class of) Σ . Consider a family of boundary regions $\mathcal{A}_\alpha \subseteq \Sigma$ specified by a real number α . For example, we can fix the shape of \mathcal{A} and let α denote the overall size; in the simplest case of $(d - 2)$ -spherically symmetric \mathcal{A} on a spatial slice of ESU_d boundary spacetime, we take $\alpha \in [0, 1]$ to be the fractional volume of the system.⁷

Now, consider the function $S(\alpha)$ defined by a *minimization* of area over all smooth bulk surfaces homologous to \mathcal{A}_α . Starting from any surface homologous to \mathcal{A}_α , we may allow it to relax, continuously decreasing the area, such as by the mean curvature flow described in Appendix A. Since the area is bounded from below, it must tend to a limit, of which one quarter in Planck units is given by $S_k(\alpha)$. This may take several values, labeled by k , depending on the initial surface chosen, corresponding to different local minima of the area functional. Generically, the surface itself will also converge to an endpoint⁸, which will be a corresponding minimal surface \mathfrak{M}_α^k . As α is varied, we expect these minima to come in a discrete set of families depending smoothly on some parameter (since \mathcal{A}_α and the spacetime are smooth), giving a set of curves in the plane of α and area. This means we have a set of functions $S_k(\alpha)$, defined on some interval of α , continuous, and smooth away from critical points of α , with $S(\alpha) = \min_k S_k(\alpha)$.

We now wish to consider the behaviour of $S(\alpha) \equiv \min_k S_k(\alpha)$. In simple spacetimes, such as pure AdS and small deformations thereof, there is only a single family $k = 1$ which covers the full range of α , so in such cases $S(\alpha)$ is manifestly smooth. However, it may happen at some point that two families, $k = 1, 2$, exchange dominance such that $S(\alpha) = S_1(\alpha) < S_2(\alpha)$ for α smaller than some critical value $\alpha_{\mathcal{X}}$, and $S(\alpha) = S_2(\alpha) < S_1(\alpha)$ for $\alpha > \alpha_{\mathcal{X}}$. At $\alpha = \alpha_{\mathcal{X}}$, $S_1 = S_2 = S$,

⁷ With the standard $SO(d - 1)$ symmetric metric on spatial sections of ESU_d , i.e., $ds_{\Sigma}^2 = d\theta^2 + \sin^2 \theta d\Omega_{d-2}^2$, we have $\alpha = \int_0^{\theta_\infty} (\sin \theta)^{d-2} d\theta / \int_0^\pi (\sin \theta)^{d-2} d\theta = \frac{1}{2} - \frac{\Gamma(\frac{d}{2})}{\sqrt{\pi} \Gamma(\frac{d-1}{2})} \cos \theta_\infty {}_2F_1\left(\frac{1}{2}, \frac{3-d}{2}, \frac{3}{2}, \cos^2 \theta_\infty\right)$ for a polar cap characterized by co-latitude θ_∞ .

⁸ For some special spacetimes, such as extremal black holes with an infinite throat, this need not be strictly true, though it does not materially affect the argument.

so $S(\alpha)$ is necessarily continuous. However, there will generically be a discontinuity in $\partial_\alpha S(\alpha)$, so $S(\alpha)$ has a kink.

We argue that this discontinuity of the derivative is as bad as it can get, and $S(\alpha)$ itself can never be discontinuous. For suppose $S(\alpha)$ has a jump discontinuity at some value $\alpha = \alpha_m$, so $S(\alpha)$ has different limits $S_- < S_+$ from the left and right respectively. Then by taking some surface with α less than, but sufficiently close to, α_m , and area sufficiently close to S_- , and deforming slightly, we expect to be able to find a surface with $\alpha > \alpha_m$, but area still less than S_+ : the variation in α may be taken as small as desired to temper the change in area. This is in contradiction with S_α being the minimum of the area functional. If desired, this deformed surface can be considered as the initial surface of a mean curvature flow. This argument precludes, for example, a discontinuity when new families of minimal surfaces appear at $\alpha = \alpha_m$, and do not exist for $\alpha < \alpha_m$. This creation can happen, and typically such families appear at α_m in pairs (we shall present an example in Fig. 8 below), but when this happens the area of the new family must exceed the area of some existing family.

The above argument indicates that within the RT prescription, the entanglement entropy of a given region \mathcal{A} should vary continuously with the parameters specifying the geometrical attributes of \mathcal{A} . In the following two sections, we will see this behaviour realized manifestly, even in situations where new families of minimal surfaces get nucleated at some intermediate α . By examining these families more closely, we will identify examples with large multiplicities of minimal surfaces.

3 Entanglement entropy in 1 + 1 dimensional CFTs

In the previous section we have argued that for the holographic entanglement entropy given by a minimization procedure, $S_{\mathcal{A}}(\alpha)$ must be continuous but need not be differentiable. The lack of differentiability is typically associated with two families of minimal surfaces exchanging dominance. Such an occurrence is not new, and good examples already exist in the literature. The simplest one occurs for the bulk spacetime being the BTZ black hole, where this point was appreciated already in the early days [1, 18, 19] and fleshed out a bit more explicitly in [20]. We quickly review this story to illustrate the contrast with our other examples in higher dimensions.

The metric for the BTZ black hole is given by⁹

$$ds^2 = -f(r) dt^2 + \frac{dr^2}{f(r)} + r^2 d\theta^2, \quad f(r) = r^2 - r_+^2. \quad (3.1)$$

It is a simple matter to find the minimal surfaces for regions $\mathcal{A} = \{\theta : |\theta| \leq \theta_\infty\}$ since these are given by spacelike geodesics. The result is best described by writing down the spatial projection of the surfaces [22]

$$\mathfrak{M}_1(\theta_\infty) : \quad \left\{ (r, \theta) : r = \gamma(\theta, \theta_\infty, r_+) \equiv r_+ \left(1 - \frac{\cosh^2(r_+ \theta)}{\cosh^2(r_+ \theta_\infty)} \right)^{-\frac{1}{2}} \right\} \quad (3.2)$$

⁹ We work in units where the AdS length $\ell_{\text{AdS}} = 1$ and also set the radius of the boundary circle R parameterized by $\theta \in [0, 2\pi]$ to unity. It is easy to restore dimensions when necessary as we illustrate later.




Fig. 1: Minimal surfaces (geodesics) in BTZ geometry with black hole size $r_+ = 1$ (left) and $r_+ = 0.2$ (right). In each panel, the thick black circle represents the AdS boundary, and the thick red circle the black hole horizon. The minimal surfaces are depicted by the thin curves, color-coded by θ_∞ , with θ_∞ varying between 0 (red) and π (purple) in increments of 0.05π . For ease of visualization we use a compactified radial coordinate $\rho = \tan^{-1} r$. The boundary region \mathcal{A} is centered at $\theta = 0$ which is plotted on top.

and is plotted in Fig. 1 for large and small black holes, for a set of $\theta_\infty \in [0, \pi]$.

As can be easily seen in Fig. 1, connected spacelike geodesics (satisfying the homology constraint) always exist for any θ_∞ and r_+ . This makes sense, since there is no reason for the geodesics to break up (in fact, spacelike geodesics can orbit the black hole arbitrarily many times before returning to the boundary, albeit at the expense of greater length). As we will see in §4, this is in stark contrast to the behaviour of co-dimension 2 surfaces in higher dimensional black hole spacetimes. Also note that in the BTZ case, another effect of the low dimensionality is that arbitrarily small black hole ($r_+ \rightarrow 0$ in (3.1)) always looks effectively large in terms of the effect it has on geodesics.

Let us now consider the proper length along these geodesics and compute the entanglement entropy. As is well known, this computation reproduces the CFT computation of Cardy-Calabrese [24] for thermal CFT on the infinite line. Naive application of the RT result for the CFT on a cylinder leads to¹⁰

$$(S_{\mathcal{A}})_{\text{naive}} = \frac{c}{3} \log \left(\frac{2r_\infty}{r_+} \sinh(r_+ \theta_\infty) \right), \quad \forall \theta_\infty \in [0, \pi] \quad (3.3)$$

where r_∞ is a radial cut-off to regulate the proper length of the geodesic and is our proxy for the

¹⁰ We use CFT central charge determined by the Brown-Henneaux analysis $c = \frac{3\ell_{\text{AdS}}}{2G_N^{(3)}}$ and note that $T = \frac{r_+}{2\pi\ell_{\text{AdS}}R}$ is the temperature of the thermal density matrix ρ_Σ .

boundary UV cut-off scale. It is then a simple matter to check that the naive result for $\delta S_{\mathcal{A}}$

$$(\delta S_{\mathcal{A}})_{\text{naive}} = \frac{c}{3} \log \left[\sinh(r_+ \theta_\infty) \operatorname{csch}(r_+ (\pi - \theta_\infty)) \right] \quad (3.4)$$

violates the Araki-Lieb inequality. For the BTZ black hole $S_{\rho_\Sigma} = \frac{\pi r_+}{2G_N^{(3)}} = \frac{2\pi^2}{3} c T R$, which is clearly bounded, while (3.4) diverges as $\theta_\infty \rightarrow 0$.

The basic point is simply the following: for regions \mathcal{A} which are sufficiently large $\theta_\infty \geq \theta_\infty^{\mathcal{X}} > \pi/2$ the holographic entanglement entropy is not computed from the connected minimal surface \mathfrak{M}_1 but rather from a shorter two-component (disconnected) one \mathfrak{M}_2 . Note that in any static black hole spacetime, the bifurcation surface is minimal by virtue of the null generator of the horizon vanishing there quadratically, forcing its extrinsic curvature to vanish. While for any \mathcal{A} there exists a connected minimal surface which satisfies the homology condition, there also exists a disconnected minimal surface, given by the union of the (connected) minimal surface for \mathcal{A}^c and the bifurcation surface of the event horizon, which likewise does the job. Thus the second family of minimal surfaces, parameterized by θ_∞ , is simply

$$\mathfrak{M}_2(\theta_\infty) : \left\{ (r, \theta) : r = r_+ \cup r = \gamma(\theta, \pi - \theta_\infty, r_+) \right\}, \quad (3.5)$$

where γ is given by (3.2) (an example of both families of surfaces is plotted in Fig. 2 explained below). Taking this fact into account we learn that the holographic entanglement entropy is given by

$$S_{\mathcal{A}}(\theta_\infty) = \min \left\{ \frac{\text{Area}(\mathfrak{M}_1)}{4G_N^{(3)}}, \frac{\text{Area}(\mathfrak{M}_2)}{4G_N^{(3)}} \right\}, \quad (3.6)$$

which then implies upon evaluating the lengths of the curves explicitly that

$$S_{\mathcal{A}} = \begin{cases} \frac{c}{3} \log \left(\frac{2r_\infty}{r_+} \sinh(r_+ \theta_\infty) \right), & \theta_\infty < \theta_\infty^{\mathcal{X}} \\ \frac{c}{3} \pi r_+ + \frac{c}{3} \log \left(\frac{2r_\infty}{r_+} \sinh(r_+ (\pi - \theta_\infty)) \right), & \theta_\infty \geq \theta_\infty^{\mathcal{X}} \end{cases} \quad (3.7)$$

where we introduce the entanglement plateau scale $\theta_\infty^{\mathcal{X}}$ to distinguish the dominant saddle point of the area functional. This can be explicitly evaluated as a function of the black hole size (or temperature) to be

$$\theta_\infty^{\mathcal{X}}(r_+) = \frac{1}{r_+} \coth^{-1} (2 \coth(\pi r_+) - 1). \quad (3.8)$$

For orientation, we plot the two sets of minimal surfaces $\mathfrak{M}_1(\alpha_{\mathcal{X}})$ and $\mathfrak{M}_2(\alpha_{\mathcal{X}})$ at the transition value $\alpha_{\mathcal{X}} = \theta_\infty^{\mathcal{X}}/\pi$ in Fig. 2 (again for two black hole sizes for ease of comparison with Fig. 1). At this value of α , \mathfrak{M}_1 (solid curve) and \mathfrak{M}_2 (dashed curves) have equal proper lengths. Note that in the high temperature (large black hole) limit, the entanglement plateau scale approaches the size of the entire system:

$$\lim_{r_+ \rightarrow \infty} \theta_\infty^{\mathcal{X}}(r_+) = \pi. \quad (3.9)$$




Fig. 2: Comparison of the two families \mathfrak{M}_1 and \mathfrak{M}_2 of the minimal surfaces (geodesics) in BTZ, with black hole sizes $r_+ = 1$ (left) and $r_+ = 0.2$ (right) as in Fig. 1, plotted at the critical value of θ_∞ where they exchange dominance, i.e. at $\theta_\infty = \theta_\infty^{\mathcal{X}}$. The connected family \mathfrak{M}_1 is represented by the solid curve (color-coded by θ_∞ as in Fig. 1), while the disconnected family \mathfrak{M}_2 is given by the two dashed curves (color-coded by $\pi - \theta_\infty$). The thick orange arc on the boundary represents the region \mathcal{A} .

Translating this into field theory quantities and using $\alpha = \frac{\theta_\infty}{\pi}$ to parameterize the fraction of the system we consider, we have

$$\alpha_{\mathcal{X}} = \frac{1}{2\pi^2 T R} \coth^{-1} (2 \coth (2\pi^2 T R) - 1) \quad (3.10)$$

where R is the radius of the boundary CFT cylinder. The main feature we want to illustrate is that for $\alpha > \alpha_{\mathcal{X}}$

$$\delta S_{\mathcal{A}} = S_{\rho_\Sigma} = \frac{2\pi^2}{3} c T R \quad \implies \quad S_{\mathcal{A}} = S_{\mathcal{A}^c} + S_{\rho_\Sigma} \quad (3.11)$$

as anticipated. Essentially for large enough boundary regions we can read off the thermal entropy by comparing directly the entanglement entropy of a region and its complement. The behaviour of $\delta S_{\mathcal{A}}$ for 1 + 1 dimensional CFTs is shown explicitly later in Fig. 5 (where we also demonstrate a similar feature in higher dimensional holographic field theories).

In deriving the relations (3.11), we have implicitly assumed $r_+ \geq 1$ (equivalently $TR \geq \frac{1}{2\pi}$), which is where the BTZ black hole is dual to the thermal density matrix for the CFT. For lower temperatures the density matrix is dual to thermal AdS geometry and the holographic computation described above should be modified. Geodesics in thermal AdS will actually give a result which says $S_{\mathcal{A}} = S_{\mathcal{A}^c}$ at leading order in the $c \rightarrow \infty$ limit.¹¹ The thermal result will only be recovered by considering $1/c$ corrections. This makes sense since for $TR < \frac{1}{2\pi}$ one is in the ‘confined’ phase of the CFT (the terminology is inherited from higher dimensions where one has an honest confinement/deconfinement transition in the large N planar gauge theory).

¹¹ This is the statement that $\mathcal{O}(1)$ corrections cannot be recovered from the classical gravity approximation.




Fig. 3: Plot of the curves $\delta S_{\mathcal{A}}/S_{\rho_{\Sigma}}$ for a Dirac fermion in 1 + 1 dimensions in the canonical ($T \neq 0$) [19] and grand canonical ($T, \mu \neq 0$) [25]. We examine the behaviour for a range of temperatures and chemical potential. The solid curves from bottom are $\beta = 4$ (red), $\beta = 2$ (blue), $\beta = 1$ (magenta) and $\beta = 0.1$ (black). The situation with the chemical potential turned on is indicated with the same colour coding with $\mu = 0.1$ represented by circles and $\mu = 0.5$ by triangles. The symmetry $\mu \leftrightarrow 1 - \mu$ is used to restrict $\mu \in [0, 1/2]$ and we see that for $\mu = 0.5$ the normalized $\delta S_{\mathcal{A}}$ is essentially the same at any temperature. This behaviour should be contrasted against the large c holographic result displayed in Fig. 5 for $d \geq 2$ thermal CFTs.

In 1 + 1 dimensions we have also the ability to compare our large c result for $S_{\mathcal{A}}$ or $\delta S_{\mathcal{A}}$ with the behaviour encountered in small central charge systems, which is known for a couple of cases. The result for the entanglement entropy of a free neutral Dirac fermion $c = 1$ was derived in [19] and was extended to a grand canonical density matrix including a $U(1)$ chemical potential in [25].¹² In either case we can use the resulting expressions (which being rather long we refrain from reproducing here) to understand the behaviour of the entanglement in the system. Using the explicit expressions one can check that the resulting answer exhibits no sharp feature. $\delta S_{\mathcal{A}}/S_{\rho_{\Sigma}}$ is a smooth monotone function of α , see Fig. 3. This is also consistent with the recent discussion of [26] who compute analogous quantities in a harmonic chain which is a gapped theory and note that sharp features can only occur in the $c \rightarrow \infty$ limit (which as we mention earlier implements the thermodynamic limit).

Note that the Dirac fermion example is so far the only case for which the entanglement entropy of a thermal CFT in finite volume is explicitly known. In general one expects that the details of the spectrum of the CFT play a crucial role in the entanglement, i.e., the answer is not a simple universal function of the central charge as the holographic result (3.7) suggests. It would

¹² We note that angular momentum chemical potential was considered in [3], which can be captured holographically using a rotating BTZ black hole in the large c limit. However the ensemble is stationary (and not static) so one needs to work with extremal surfaces as emphasized there.

be interesting to see if one can use the technology of [9] to argue directly from a CFT analysis for an entanglement plateau relation like (3.11) in the asymptotic large c limit.

The behaviour of $S_{\mathcal{A}}(\alpha)$, not surprisingly, is in fact rather similar to the behaviour of the entropy or free energy in the thermal ensemble (a point we will revisit in §7). In both cases there are multiple saddle points, which exchange dominance. For the thermal density matrix the saddles are the thermal AdS geometry and the BTZ black hole both of which exist for the entire range of the dimensionless parameter TR . For the holographic entanglement entropy $S_{\mathcal{A}}$ we have again two distinct saddles available for the entire range of parameters: the connected and disconnected surfaces \mathfrak{M}_1 and \mathfrak{M}_2 exist for all values of $\alpha \in [0, 1]$. For a sufficiently large region $\alpha \geq \alpha_{\mathcal{X}}$, however, it is the disconnected surface that dominates and results in the entanglement plateau. We will soon see that existence of both families over the entire range of α is peculiar to three dimensions and the situation is much more intricate in higher dimensions.

4 Entanglement entropy in higher dimensional field theories

In the previous section we saw that already in the 3-dimensional bulk spacetime corresponding to a thermal state in the dual CFT, we have multiple families of minimal surfaces \mathfrak{M}_i anchored at the boundary entangling surface $\partial\mathcal{A}$, for arbitrary \mathcal{A} . Let us now examine the analogous situation in higher dimensions. We will consider the Schwarzschild-AdS $_{d+1}$ bulk spacetime ($\ell_{\text{AdS}} = 1$),

$$ds^2 = -f(r) dt^2 + \frac{dr^2}{f(r)} + r^2 (d\theta^2 + \sin^2 \theta d\Omega_{d-2}^2), \quad f(r) = r^2 + 1 - \frac{r_+^{d-2}(r_+^2 + 1)}{r^{d-2}}. \quad (4.1)$$

Large black holes $r_+ \geq 1$ describe the thermal state on ESU (e.g. of $\mathcal{N} = 4$ SYM on $\mathbf{S}^3 \times \mathbb{R}$ in the best-understood case of $d = 4$), but we will consider black holes of any size for generality. For simplicity, we will take the boundary region \mathcal{A} to be a disk centered at the ‘north pole’ $\theta = 0$ with radius θ_∞ , and consider only surfaces which maintain the residual $SO(d-1)$ spherical symmetry and remain at constant t . This effectively reduces the problem to a 2-dimensional one: we can specify any such minimal surface as a curve in the (r, θ) plane.¹³

The equations of motion for the minimal surface are obtained from the Nambu-Goto action for minimizing the area of the surface. Using an auxiliary parameter s we have the Lagrangian for the system

$$\mathcal{L} = (r \sin \theta)^{d-2} \sqrt{\frac{1}{f(r)} \left(\frac{dr}{ds}\right)^2 + r^2 \left(\frac{d\theta}{ds}\right)^2}. \quad (4.2)$$

The Euler-Lagrange equations for $r(s)$ and $\theta(s)$ are equivalent due to the reparameterization invariance, giving a second order ODE for $r(s)$ and $\theta(s)$. The second equation comes from choice of parameter.¹⁴

¹³ For plotting purposes, we will consider a compactified radial coordinate $\rho = \tan^{-1} r$ as in §3, and double up the $\theta \in [0, \pi]$ coordinate to $\theta \in [-\pi, \pi]$ so that all curves will be reflection-symmetric. The black hole will then be represented by a disk of radius $\tan^{-1} r_+$ centered at the origin.

¹⁴ One might be tempted to use θ itself as a parameter, but it turns out that θ is not monotonic in higher

The topology of the problem constrains any surface to pass through a pole of the sphere, so to classify all connected surfaces we may start integration (w.l.o.g.) at the North pole. Requiring smoothness of \mathfrak{M} there, we obtain a one-parameter family of minimal surfaces, specified by the ‘initial value’ for the radial coordinate at the north pole, $r_0 \equiv r(\theta = 0)$.¹⁵ Given r_0 , we find the minimal surface and from this read off the latitude θ_∞ which it is anchored on. From this we will classify the boundary region size by α , the proportion of the area contained in the boundary region homologous to the surface: $\alpha = \frac{\text{Vol}(\mathcal{A})}{\text{Vol}(\mathbf{S}^{d-1})}$ which tends to 0, 1 as $\theta_\infty \rightarrow 0, \pi$ respectively.

4.1 Entanglement plateaux in $d > 2$

Considering α as a function of initial radius r_0 , the situation is rather different from the $d = 2$ case. In $d > 2$, α reaches a maximum value of $\alpha_m < 1$: for sufficiently large regions, there are no connected minimal surfaces obeying the homology constraint! This can be intuited from analogous behaviour in the classic ‘soap bubble’ problem in flat space, of finding a minimal surface between two circular rings. To reduce the area, there is a tendency to shrink the radius of the tube, counteracted by the constraint of ending on the rings. But when the ratio of ring radius to separation is sufficiently small, the rings do not hold the surface up enough to prevent the tube radius from shrinking to zero, and the surface separates into two disconnected parts. This reasoning carries over directly to our set-up, and the process of surfaces splitting into disconnected pieces can be seen explicitly in the animations of our numerical simulations of mean curvature flow, to be found on the arXiv as [ancillary files](#) for the submission. The area cost of having a wide tube is greater in higher dimensions, which leads one to expect that α_m should be smaller, reflecting the greater tendency of the surfaces to split up.¹⁶ The physical point is that this leaves no option but to consider the disconnected surfaces, to which we now turn. In the next section, we offer a very different geometrical justification of why connected extremal surfaces homologous to \mathcal{A} cannot exist beyond a certain θ_∞ , which is based on causality in the full Lorentzian geometry.

So far, we have a one-parameter family of connected surfaces \mathfrak{M}_1 . In addition to this, we must consider the ‘disconnected’ family \mathfrak{M}_2 of surfaces with two connected components: the first, anchored to the boundary, is a reflection of a connected surface already considered, passing through the South pole. Due to the nontrivial $\mathbf{S}^{d-1} \times \mathbb{R}$ topology of the static slice of the geometry, this alone fails to satisfy the homology constraint, so it is supplemented by a second piece, the bifurcation sphere on the event horizon of the black hole. To identify the correct entangling surface

dimensions. We choose s such that the evaluation of the area integral on-shell reduces to $\int ds (r \sin \theta)^{d-2}$, for good behaviour numerically. In particular, the results for the regularized area of the surface are less error-prone since r increases exponentially in s as the boundary is approached.

¹⁵ Apart from ODE methods, the numerical construction of the surfaces is also done using a mean curvature flow. The reader interested in the details is encouraged to consult Appendix A where we outline the necessary mathematical technology and the algorithm used for obtaining the surfaces.

¹⁶ This behaviour is in fact reminiscent of similar observation in gravitational context of the Gregory-Laflamme instability of higher-dimensional black strings (or branes), which are likewise more prone to fragmenting with increasing dimension [27, 28].




Fig. 4: Minimal surfaces for $r_+ = 1$ (left) and $r_+ = 0.2$ (right) black holes in Schwarzschild-AdS in $3 + 1$ (top) and $4 + 1$ (bottom) dimensions, analogous to the plots of Fig. 1. In each panel, the thick black circle represents the AdS boundary, and the thin red circle the black hole horizon. The boundary region \mathcal{A} is centered at $\theta = 0$ (North Pole) which is plotted on top. Further plots with other black hole sizes and higher dimensions can be found as [ancillary files](#) with the submission.

determining $S_{\mathcal{A}}(\alpha)$, we must compare the areas of these two families; they exchange dominance at $\alpha_{\mathcal{X}} < \alpha_m$, as is inevitable from continuity. This means that at this value $\alpha_{\mathcal{X}}$, there are two surfaces of equal area, and as we vary α through $\alpha_{\mathcal{X}}$, the surfaces jump and the entanglement entropy as a function of region size has a kink. In Fig. 4 we show the surfaces which determine $S_{\mathcal{A}}(\alpha)$ for various black hole sizes and dimensions. Those passing ‘above’ the black hole belong to the family \mathfrak{M}_1 , and those passing ‘below’ should be supplemented by the horizon, and belong to \mathfrak{M}_2 . Comparing the panels horizontally, we see that the smaller the black hole, the smaller $\alpha_{\mathcal{X}}$ gets; indeed, as $r_+ \rightarrow 0$, $\alpha_{\mathcal{X}} \rightarrow 1/2$, whereas as $r_+ \rightarrow \infty$, $\alpha_{\mathcal{X}} \rightarrow 1$; we plot the actual curve




Fig. 5: Entanglement entropy plateaux for Schwarzschild-AdS black holes dual to thermal field theories. (a). Behaviour in Schwarzschild-AdS₅ for varying r_+ . From right to left we have $r_+ = \frac{1}{4}$ (blue), $r_+ = \frac{1}{2}$ (purple), $r_+ = 1$ (yellow), $r_+ = 2$ (green) and $r_+ = 4$ (blue). (b). Dimension dependence of the plateaux, at the Hawking-Page transition point $RT = \frac{d-1}{2\pi}$ for $d = 2, 3, 4, 5, 6$ (red, yellow, green, blue, purple respectively). Note that the horizon size is related to the temperature via $r_+ = \frac{2\pi}{d} \left(RT + \sqrt{(RT)^2 - \frac{d(d-2)}{4\pi^2}} \right)$.

$\alpha_{\mathcal{X}}(r_+)$ in §5. On the other hand, comparing the panels vertically, we see that as d increases, the surfaces get less affected by the black hole until close to the horizon. This is easy to understand from the simple fact that gravity falls off faster in higher dimensions, and chimes perfectly with the intuition recently explained in [28]. Turning to common characteristics, one universal feature of all cases is that none of the minimal surfaces (anchored on the boundary) can penetrate past the event horizon. This is in fact true for any static geometry with a horizon [29].¹⁷

It is evident from the symmetry of our setup that in the regime $\alpha \geq \alpha_{\mathcal{X}}$, we have the identity

$$S_{\mathcal{A}} = S_{\mathcal{A}^c} + S_{\rho_{\Sigma}} \quad (4.3)$$

so that we manifestly saturate the Araki-Lieb inequality. This is identical to the situation described in §3. Indeed, a plot of $\delta S_{\mathcal{A}}$ as a function of α for various dimensions explicitly reveals the entanglement plateaux as illustrated in Fig. 5 where we include the result for $d = 2$ CFTs as well for completeness. In Fig. 5(a) we show the behaviour for a fixed dimension $d = 4$ as we vary the temperature by changing the black hole size, whereas in Fig. 5(b) we display the behaviour for different dimensions at fixed $TR = \frac{d-1}{2\pi}$. We see that the entanglement plateau ends by a kink (e.g. in $d = 2$ the behaviour of $\delta S_{\mathcal{A}}/S_{\rho_{\Sigma}}$ can be obtained from (3.7)), the extent of which gets smaller with increasing dimension. To be sure, the precise details depend on the nature of the comparison; we have found it reasonable to fix the temperature (in units of the boundary sphere) and compare the behaviour across dimensions. In general the plateau width decreases as a function of T for a

¹⁷ More specifically, [29] showed that for asymptotically AdS spacetimes with planar symmetry the deepest reach (i.e. the turning point, furthest away from the boundary) of any extremal surface which is fully anchored on (a single) AdS boundary cannot occur inside a black hole. On the other hand, in time-dependent spacetimes, there do exist extremal surfaces anchored on the boundary which penetrate past the horizon [12, 29–33]. For this reason, entanglement entropy is cognizant of some physics inside the horizon.




Fig. 6: Continuation of family \mathfrak{M}_1 of minimal surfaces in Schwarzschild-AdS₅ for $r_+ = 1$. As the turning point r_0 approaches the horizon, we see an intricate pattern of the minimal surfaces folding back onto themselves; solutions with n -folds are denoted $\mathfrak{M}_{1,(n)}$. We illustrate $\mathfrak{M}_{1,(0)}$ (red, yellow) and $\mathfrak{M}_{1,(1)}$ (blue, purple) surfaces for two specific choices of α ; these two families terminate at $\alpha = \alpha_m$ (green).

given dimension (the behaviour as a function of d depends on the precise value of RT ; see [ancillary files](#) for $RT = \frac{2}{3}$).

4.2 Sub-dominant saddles: folds in minimal surfaces

Although we now have the curve $S_A(\alpha)$ for the full range of α , it is interesting to ask what happens to the connected minimal surfaces (and hence the disconnected ones as well, by the reflection symmetry) in their sub-dominant regime. We know that a continuation of these families must exist, since the deepest radius r_0 did not cover the full range down to r_+ .

As we decrease r_0 from the critical value $r_0(\alpha_{\mathcal{X}})$, the endpoint of the surface α increases beyond $\alpha_{\mathcal{X}}$ as expected, but as previously noted, reaches a maximum at $\alpha = \alpha_m$, beyond which there simply are no connected surfaces at all. Instead, as one decreases r_0 further, the endpoints α start to recede to lower values again, even though the ‘neck’ of the surface \mathfrak{M}_1 near $\theta = \pi$ keeps closing off. This behaviour is illustrated in Fig. 6. Bringing r_0 still closer to the horizon reveals another surprise: the pattern repeats itself. Eventually α reaches a minimal value (which happens to be very close to $1 - \alpha_m$) and turns around again. In other words, now the corresponding minimal




Fig. 7: Near-horizon analysis of folds in the minimal surfaces for $r_+ = 1$ (left) and $r_+ = 0.2$ (right) black holes in Schwarzschild-AdS₄. Here $x_0 = \frac{1}{2} \log(r_0 - r_+)$ measures the closest approach of the surface to the horizon.

surface has two necks, one near $\theta = \pi$ and the other near $\theta = 0$.

Careful examination reveals that the closer the minimal radius r_0 is to the horizon, the more intricate the surface. The endpoints θ_∞ are restricted to lie sufficiently close to the equator $\theta = \frac{\pi}{2}$, in particular, $1 - \alpha_m \lesssim \alpha \lesssim \alpha_m$, as manifested in Fig. 7, where we plot α as a function of rescaled coordinate $x_0 = \frac{1}{2} \log(r_0 - r_+)$ indicating the proximity of r_0 to the horizon (cf Appendix B).

In fact, the pattern of minimal surfaces exhibits a discrete self-similarity. The minimal surfaces wrap the horizon multiple times, the distance from the horizon of each wrapping parametrically separated from the next. We denote the solution with n -folds around the horizon as $\mathfrak{M}_{1,(n)}$. The connected surfaces described earlier in Fig. 4 are $\mathfrak{M}_{1,(0)}$ with this refinement in the notation. The existence of these solutions and their properties can be analytically understood using a local analysis around the turning points, and near the horizon away from the turning points, as we illustrate in Appendix B. Further, this analysis is generic for any spherically symmetric non-extremal horizon, so there is nothing particularly special about Schwarzschild-AdS in this regard.

Of course, such surfaces which fold around the black hole multiple times have larger area (approximately by the black hole area times the number of foldings) than their simpler cousins, so these surfaces are not directly relevant for the entanglement entropy. This is already visible explicitly in Fig. 8, where we plot the actual areas of the minimal surfaces, as a function of α . For $\alpha > \alpha_\chi$ the disconnected family \mathfrak{M}_2 takes over, a fact we have already used to demonstrate the entanglement plateau in Fig. 5. In the regime $\alpha > \alpha_\chi$ the connected surfaces have greater area, and indeed the area continues to grow as the surface folds over about the horizon. It is intriguing to note that the difference $\text{Area}(\mathfrak{M}_{1,(n)}) - \text{Area}(\mathfrak{M}_{1,(n-2)}) \approx 2 S_{\rho_\Sigma}$ increasingly more accurately as n gets large. While the sub-dominant connected surfaces $\mathfrak{M}_{1,(n)}$ for discussed above are irrelevant for the boundary entanglement for $\alpha > \alpha_\chi$ it is nevertheless curious that such a simple geometry as Schwarzschild-AdS allows for such a rich structure of minimal surfaces!

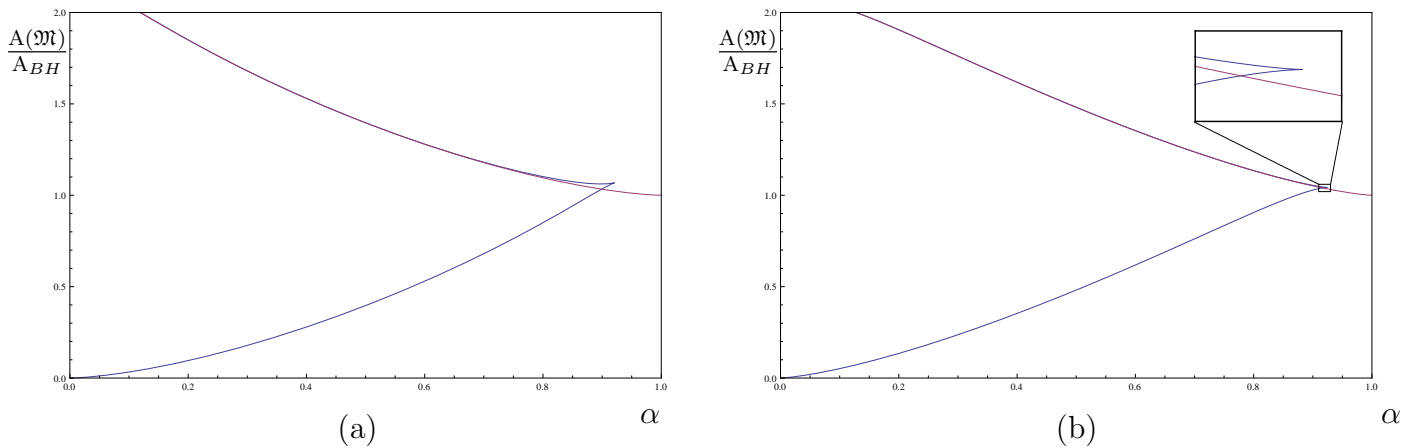


Fig. 8: Regulated of minimal surfaces (normalized in units of black hole area), for (a) Schwarzschild-AdS₄ ($r_+ = 1$), and (b) Schwarzschild-AdS₅ ($r_+ = 1$). The blue curve is the family \mathfrak{M}_1 , and the purple curve the disconnected family \mathfrak{M}_2 . The area of the connected minimal surfaces grows without bound, since we encounter multiply-folded surfaces $\mathfrak{M}_{1,(n)}$ which for $n > 1$ which wrap the horizon multiple times. The swallow-tail behaviour characteristic of first order phase transitions is clearly visible in the plots.

5 Causality and holographic entanglement

In §4 we have examined the families of minimal surfaces in static thermal density matrices for holographic theories which compute the entanglement entropy. The main surprise was that beyond a certain size of \mathcal{A} , namely for $\alpha > \alpha_m$, there are no connected surfaces homologous to \mathcal{A} , which automatically forces the holographic entanglement entropy to plateau in the sense described earlier. Above, we demonstrated this result by explicitly constructing minimal surfaces in the bulk Schwarzschild-AdS _{$d+1$} spacetime; the robustness of numerical construction rests on explicitly allowing surfaces to relax to minimal ones using a mean curvature flow as explained in Appendix A. However, the absence of connected minimal surfaces is in fact necessitated by a certain relational property of minimal surfaces and causal wedges. The latter is a construct that crucially involves the Lorentzian structure since it has to do with causal relations between points. We now proceed to explain the basic physics behind this connection, leaving the technical details for a separate discussion [23].

Given a boundary region $\mathcal{A} \subset \Sigma$, one can construct a causal wedge (which, following [22], we denote by $\blacklozenge_{\mathcal{A}}$) associated with it in the bulk. The construction involves first identifying the boundary domain of dependence $\diamond_{\mathcal{A}}$ of \mathcal{A} , which for the rigid boundary geometry is simply the region where one can reliably Cauchy-evolve initial data laid down on \mathcal{A} . If one has full knowledge of $\rho_{\mathcal{A}}$ then one can compute correlation functions of all local operators inserted in $\diamond_{\mathcal{A}}$; this is just a basic statement of causality in relativistic quantum field theories. We then construct the causal wedge $\blacklozenge_{\mathcal{A}}$ in the bulk, defined as the set of points in \mathcal{M} which can communicate with and simultaneously be communicated from $\diamond_{\mathcal{A}}$; in other words, through which there exists a causal curve which starts and ends in $\diamond_{\mathcal{A}}$. The reader interested in the precise definitions is invited to

consult [22].

Since the causal wedge is constructed purely based on where causal curves reach into the bulk, one might naively imagine that for a simply-connected boundary region \mathcal{A} the causal wedge has trivial topology. This turns out not to be the case [23]. While the causal wedge for a simply connected region has to be simply connected (a consequence of topological censorship), it can have ‘holes’. This implies that $\blacklozenge_{\mathcal{A}}$ possesses non-contractible $d-2$ spheres on its boundary. Intuitively, this can happen because of geodesic trapping, as we now explain. The boundary of the causal wedge is generated by null geodesics. Bulk spacetimes which possess null circular orbits (arising from a sufficiently strong gravitational potential) cut off some of these generators, while allowing some others to fly-by and thereby create a hole in the wedge. Intuitively, the play-off is always between gravitational attraction and centrifugal repulsion and for a wide family of geometries one is able to engineer these so as to obtain the desired effect. Happily for us, the simplest class of examples where such a feature can be exhibited explicitly are the Schwarzschild-AdS $_{d+1}$ geometries, wherein for large enough region \mathcal{A} (which we denote by $\alpha > \alpha_{\blacklozenge}$), the causal wedge $\blacklozenge_{\mathcal{A}}$ exhibits a hole [23].

The presence of non-trivial topology has an important implication for the minimal surfaces \mathfrak{M} . This is at first sight surprising, as one might have naively thought that since minimal surfaces are defined in the Riemannian section (constant-time slice), they ought to be utterly ignorant of any causal argument which is intrinsically Lorentzian. This naive reasoning is however too quick: One can show that all extremal surfaces $\mathfrak{E}_{\mathcal{A}}$ (of which the minimal surfaces $\mathfrak{M}_{\mathcal{A}}$ considered herein are a special case) must lie outside (or at best on the boundary of) the causal wedge [22, 23, 34]. Here by ‘outside’ we mean ‘not within’ – in other words an extremal surface $\mathfrak{E}_{\mathcal{A}}$ must penetrate deeper into the bulk than the corresponding causal wedge $\blacklozenge_{\mathcal{A}}$. The result can be intuitively understood by realizing that a minimal surface wants to minimize its area and thus wants to spend much of its existence deep in the interior where it can minimize the red-shift factor of AdS (which augments the area element of the directions along the boundary). This effect is less dominant for lower-dimensional surfaces, and null geodesics can additionally offset the large spatial distances by large temporal ones.¹⁸ Hence generically the causal wedge boundary is nestled between the extremal surface and the AdS boundary. As an aside, this statement forms a crucial ingredient in the arguments of [35] who argue that the holographic dual of the reduced density matrix $\rho_{\mathcal{A}}$ must comprise of a bulk region that is larger than the casual wedge $\blacklozenge_{\mathcal{A}}$.¹⁹

We are now in a position to explain the behaviour found in §4. For $\alpha > \alpha_{\blacklozenge}$ the ability of the causal wedge to develop holes implies that we should anticipate a corresponding change in the

¹⁸ Note that in comparison to pure AdS geometry, both null geodesics and extremal surfaces are nevertheless pushed towards the boundary by the gravitational potential well of a deformed bulk spacetime. This in effect means that both entanglement entropy and causal holographic information defined in [22] grow with positive mass deformations of the spacetime.

¹⁹ *Caveat:* The statements made above should be viewed with suitable caution, as they require some work to be established rigorously. For further discussion and a proof of the nesting property of causal wedges and extremal surfaces (modulo some technical assumptions), we refer the reader to [23] (see also [34] for a related discussion in causally trivial spacetimes).

minimal surface. Indeed $\mathfrak{M}_{\mathcal{A}}$ cannot pass into the hole in $\blacklozenge_{\mathcal{A}}$ while remaining connected to the boundary, because of the nesting property: in order to do so, it would have to pass through the causal wedge. This means that it either lies completely outside the causal wedge or is contained entirely within the hole! One fact which is somewhat obvious by causality is that the hole in the causal wedge must lie outside the black hole event horizon. For static black hole spacetimes we have argued that the bifurcation surface is a candidate extremal surface in §3. So in the presence of the hole in the wedge we can have a minimal surface component on the far side of the black hole (hence outside $\blacklozenge_{\mathcal{A}}$) and another component being simply the bifurcation surface. In fact, as in the BTZ discussion, the homology constraint imposed upon the RT prescription by the Araki-Lieb inequality necessitates both components. Hence the only surfaces that are both minimal and homologous to \mathcal{A} are the disconnected surfaces \mathfrak{M}_2 described in §4.




Fig. 9: The critical values of α encountered in our discussion as a function of r_+ . We plot the largest fraction of the boundary region characterized by α for which (i) a connected minimal surface exists α_m (blue), (ii) for which disconnected/connected exchange dominance α_{χ} (purple) and, (iii) the critical α for which the causal wedge develops holes α_{\blacklozenge} (olive) in the Schwarzschild-AdS₅ spacetime.

We demonstrate explicitly that $\alpha_{\chi} \leq \alpha_m \leq \alpha_{\blacklozenge}$ as one anticipates for first order transitions for Schwarzschild-AdS₅ in Fig. 9. As the figure indicates, the causal wedge pinch off happens at $\alpha = \alpha_{\blacklozenge}$ which is significantly larger than the bound α_m on existence of corresponding extremal surfaces. In other words, the connected minimal surface ceases to exist long before this is necessitated by the topology of the causal wedge. While this may seem to somewhat weaken our arguments regarding the utility of the causal wedge, we should note that the pinch-off value $\alpha = \alpha_{\blacklozenge}$ is a rather weak bound from the causal wedge standpoint as well: already before the causal wedge pinches off, its geometry precludes the requisite minimal surface. This is because prior to the pinch-off the opening becomes too long and narrow to admit any extremal surface. To support the latter, the neck would have to accommodate a sufficient ‘flare-out’ shape, as explained in §4.1, which occurs for θ_{∞} smaller by $\sim \mathcal{O}(\rho_0 - \rho_+)$, where ρ_0 is the ρ -radius of null circular orbit and ρ_+ gives the horizon size. Hence a more sophisticated analysis which takes into account the shape of the causal

wedge apart from its topology would provide a much stronger bound on α_m . It would be interesting to examine in detail precisely how close we can get to α_m , but we leave this for future exploration.

6 Minimal versus extremal surfaces and the homology constraint

So far, we have examined the properties of the entanglement entropy as given by the RT proposal involving minimal surfaces (which is a well-defined problem if we can look for our surfaces in a Riemannian geometry). This allowed us to argue in §2 that the entanglement entropy for a certain region \mathcal{A} , measured in a fixed total density matrix ρ_Σ , is continuous as a function of parameters α describing the region \mathcal{A} . On the other hand, the general prescription for holographic entanglement entropy must be formulated in a fully covariant fashion. In other words, in absence of a geometrically-preferred bulk spacelike slice (which for static cases we could take to be a constant- t slice where $\frac{\partial}{\partial t}$ is the timelike Killing field), we need to have a prescription which does not rely on a specific coordinate choice. This motivated HRT [3] to propose the extremal surface prescription (1.2), which is fully covariant and well-defined for arbitrary time-dependent asymptotically-AdS bulk geometries. It has henceforth been assumed that the two prescriptions, RT and HRT, are equivalent for static geometries. In this section we wish to revisit this assumption, in conjunction with the closely-related issue of the homology requirement.

Let us start by specifying the prescriptions more explicitly, for completeness also adding a recent ‘maximin’ reformulation of HRT by Wall [34]. In each formulation, the surface is required to be anchored on $\partial\mathcal{A}$, satisfy the homology constraint, and in case of multiple surfaces satisfying these criteria, be the minimal-area one. The differences in the prescriptions enter at the level of constructing the requisite surfaces, and can be summarized as follows:

- RT [1, 2]: minimal surface on constant- t spacelike slice Σ_t^{bulk}
- HRT [3]: extremal surface in the bulk \mathcal{M}
- Wall [34]: minimal surface on any achronal bulk slice $\tilde{\Sigma}$, maximized over all possible $\tilde{\Sigma}$.

While RT is restricted to static spacetimes, Wall’s prescription is formulated in causally trivial bulk geometries. In this context, [34] proves equivalence between the maximin and HRT prescriptions assuming the null energy condition. Although the maximin construction is useful for some purposes (for example, it allows [34] to argue the existence of such surfaces and prove strong subadditivity in the time-dependent context), it is conceptually more complicated since the requisite surface is obtained by a two-step procedure of first minimizing the area on some achronal slice $\tilde{\Sigma}$, and then maximizing the area with respect to varying $\tilde{\Sigma}$. Moreover, here we wish to consider causally non-trivial spacetimes, so we will henceforth restrict attention to the RT and HRT proposals.

In a globally static geometry, RT and HRT proposals are indeed equivalent, since any extremal surface anchored at constant t on the boundary must coincide with the minimal surface on Σ_t^{bulk} ,

cf. [3].²⁰ However, in a static but not globally static geometry (i.e. when there is a global Killing field which is timelike near the AdS boundary but does not necessarily remain timelike everywhere in the bulk), the situation can appear more subtle.

To illustrate the point, let us consider the eternal non-extremal Reissner-Nordstrom-AdS geometry, corresponding to a static charged black hole. It has a metric of the form (4.1), but with $f(r)$ having two positive real roots at $r = r_{\pm}$ with r_+ corresponding to the outer (or event) horizon and r_- the inner (or Cauchy) horizon with $0 < r_- < r_+$.²¹ The causal structure is indicated in the Penrose diagram in Fig. 10. Apart from the bifurcation surface of the event horizon at $r = r_+$ which has zero extrinsic curvature (and therefore is a compact extremal surface), there is an analogous compact extremal surface at the inner horizon $r = r_-$. At both surfaces the $\frac{\partial}{\partial t}$ Killing field vanishes (and its norm decreases in spacelike directions and increases in timelike directions), but whereas the area of the surface is minimized to spatial deformations and maximized to temporal ones in the case of event horizon, it is the other way around for the inner horizon.

Outside the black hole, the geometry is qualitatively similar to that of Schwarzschild-AdS, in the sense that e.g. causal wedges for large enough boundary regions \mathcal{A} will have holes. This will in turn preclude the existence of connected extremal surfaces anchored on $\partial\mathcal{A}$ and homologous to \mathcal{A} in this regime, by the argument of §5. In such a regime, the RT prescription then instructs us to take the surface homologous to \mathcal{A}^c , along with the compact surface wrapping the event horizon. On the other hand, in the HRT prescription, there would a-priori be nothing to prevent us from taking the latter component to wrap the inner horizon instead: its area is manifestly smaller than that of the event horizon, and it still satisfies the (naive) homology constraint, as indicated pictorially by the green curve in right panel of Fig. 10. In particular, there exists a smooth co-dimension one surface whose only boundaries are the inner horizon, the extremal surface homologous to \mathcal{A}^c , and the region \mathcal{A} on the boundary.

Note that if this were indeed the correct prescription, then we would find that, instead of saturating the Araki-Lieb inequality as in (4.3), we would only satisfy it: $\delta S_{\mathcal{A}}$ would still plateau as a function of θ_{∞} , but at a value which is lower²² than the expected value $S_{\rho_{\Sigma}}$:

$$\delta S_{\mathcal{A}}(\theta_{\infty} \geq \theta_{\infty}^{\mathcal{X}}) = \frac{\Omega_{d-1} r_-^{d-1}}{4 G_N^{(d+1)}} = \left(\frac{r_-}{r_+}\right)^{d-1} S_{\rho_{\Sigma}}. \quad (6.1)$$

Although this result is consistent with the Araki-Lieb inequality, it is nevertheless at odds with

²⁰ It is worth commenting that the proof presented in [3] shows that the HRT proposal reduces to the minimal surface RT proposal in static geometries. The question of whether there are extremal surfaces not localized on a constant time slice and having smaller area than the minimal surface is, as far as we are aware, open. One might argue based on Euclidean continuation of static geometries that such surfaces do not play any role in determining $S_{\mathcal{A}}$; we thank Tadashi Takayanagi for useful discussions on this issue.

²¹ For example, in 4+1 bulk dimensions, $f(r) = r^2 + 1 - \frac{M}{r^2} + \frac{Q^2}{r^4} = \left(1 - \frac{r_+^2}{r^2}\right) \left(1 - \frac{r_-^2}{r^2}\right) (r^2 + 1 + r_+^2 + r_-^2)$.

²² The special case of extremal RN-AdS black hole with $r_- = r_+$ is somewhat more subtle, since there is no bifurcation surface (instead the spacial geometry exhibits an infinite throat). We think this is a feature rather than a bug, indicative of being at strictly zero temperature; however the HRT prescription can be applied in the same limiting fashion as is commonly done with e.g. the Wald entropy of extremal black holes.




Fig. 10: A sketch of the Penrose diagram for non-extremal Reissner-Nordstrom-AdS black hole (the figure repeats in the vertical direction, indicated by the ellipsis). The AdS boundaries are indicated by the black vertical lines, the curvature singularities by purple vertical wiggly curves, the horizons by blue diagonal dashed line, the projection of the boundary region \mathcal{A} by the orange dot, and the projection of the extremal surface by the blue horizontal line. The two panels distinguish the RT (left) and naive HRT (right) prescriptions in case of disconnected surfaces: in the former the disconnected surface lies at the event horizon (red dot) with the homology region \mathcal{R} indicated by the red dotted line. In the HRT case, the disconnected surface is at the Cauchy horizon (green dot), and corresponding \mathcal{R} then interpolates between this surface and its other boundaries as indicated by the green dotted curve.

the CFT expectation: for nearly-neutral black holes where $r_- \ll r_+$ we should be close to the thermal value rather than parametrically separated from it!

This observation suggests that we need some modification to the homology constraint specification in (1.2): the mere presence of *some* smooth surface \mathcal{R} with the requisite boundaries does not seem to suffice. One natural restriction, which indeed has already been employed in [34], is to require that \mathcal{R} be everywhere spacelike.²³ With this additional restriction, the above example would be safely invalidated, since \mathcal{R} cannot reach from the Cauchy horizon bifurcation surface $r = r_-$ to the AdS boundary $r = \infty$ while remaining spacelike everywhere, as evident from Fig. 10. The only other compact extremal surface which is spacelike-separated from the boundary region \mathcal{A} and the extremal surface homologous to \mathcal{A}^c is the event horizon bifurcation surface, which recovers the thermal answer (4.3), consistently with our expectations.

However, while the spacelike restriction on the homology constraint recovers the thermal answer for the global eternal charged black hole, there are other geometries where this does not suffice. As our second exhibit, consider a Schwarzschild-AdS ‘bag of gold’ geometry, discussed in

²³ Alternately, we could require that no component of the extremal surface is allowed to lie in the causal future of \mathcal{A} . We thank Matt Headrick and Don Marolf for useful discussions on this point.




Fig. 11: Sketch of the Schwarzschild-AdS ‘bag of gold’ embedding diagram (left) and Penrose diagram (right). Here the right asymptotic boundary of eternal Schwarzschild-AdS geometry is cut-off by a static shell (brown dotted curve), beyond which the spacetime caps off through a smooth origin as indicated. The embedding diagram depicts event horizon bifurcation surface (red curve) and the boundary region \mathcal{A} (orange curve), at the endpoints of which is anchored the connected part of the extremal surface (blue curve). The Penrose diagram has the same conventions as in Fig. 10, but now the homology surface \mathcal{R} (dotted green line) goes all the way around the tip.

e.g. [21, 36]. This has causal structure and a spatial embedding geometry as sketched in Fig. 11. The right asymptotic region, as well as interior of the black hole and white hole, are the same as in the eternal Schwarzschild-AdS geometry, but the left asymptotic region is modified by a presence of a shell whose interior has a smooth origin. Moreover, one can fine-tune the shell’s trajectory such that it remains static – so the entire spacetime admits a Killing field $\frac{\partial}{\partial t}$. In the CFT dual, such static geometry describes some equilibrium mixed (though not precisely thermal) density matrix.

Let us once again examine the entanglement entropy of a sufficiently large region \mathcal{A} , for which the causal wedge has a hole, so that no connected extremal surface anchored on $\partial\mathcal{A}$ can pass on \mathcal{A} ’s side of the black hole. In the global eternal Schwarzschild-AdS spacetime, the homology constraint would then have forced us to take the connected surface on \mathcal{A}^c ’s side of the black hole, along with the bifurcation surface on the event horizon. In the present case however, while the bifurcation surface is still an extremal surface (cf. the red curve in Fig. 11), its inclusion in the entanglement entropy computation is no longer required by the homology constraint (even including the spacelike restriction): there exists a smooth spacelike co-dimension 1 region \mathcal{R} (indicated by the dotted green line in the Penrose diagram of Fig. 11), wrapping the bag of gold at the given instant in time, whose only boundaries are \mathcal{A} and the extremal surface for \mathcal{A}^c (denoted by the blue curve in Fig. 11). In other words, in this geometry,

$$S_{\mathcal{A}} = S_{\mathcal{A}^c} \quad \Longrightarrow \quad \delta S_{\mathcal{A}} = 0 \quad \forall \theta_{\infty}. \quad (6.2)$$




Fig. 12: Sketch of the Schwarzschild-AdS joined to Schwarzschild-dS geometry across a shell: embedding diagram (left) and Penrose diagram (right), with same conventions as in Fig. 11. In contrast to the capped off geometry, however, here we do have a smaller (but non-zero) area extremal surface (green circle/dot) which is required by the homology constraint.

We stress that while the CFT is not in the precisely thermal state, it is certainly not in a pure state either [21], so a relation of the form (6.2) is wholly unexpected.²⁴

In comparing the RT and HRT prescriptions for the Schwarzschild-AdS bag of gold case, we encounter a slight ambiguity in the RT prescription, namely in what is really meant by a ‘constant t slice’: strictly-speaking, the static coordinate patch ends at the horizon (which also follows from thinking about the Euclidean section), so the RT prescription should not see any difference between this and the eternal Schwarzschild-AdS geometry – hence predicting that $\delta S_{\mathcal{A}} = S_{\rho_{\Sigma}}$ for large regions. On the other hand, if one were instructed to take a full (geodesically complete) time-symmetric slice through the global geometry, one would reproduce the HRT result (6.2). We also note that one could consider intermediate cases, with any $\delta S_{\mathcal{A}} \leq S_{\rho_{\Sigma}}$, by putting a Schwarzschild-dS geometry with smaller event horizon $\tilde{r}_+ < r_+$ on the left of the shell, as indicated in Fig. 12. The corresponding embedding diagram for a time-reflection-symmetric slice would not cap off, necessitating the inclusion of a bifurcation surface for this Schwarzschild-dS geometry in the entanglement entropy computation, in order to satisfy the homology constraint.

We have presented several examples of static (though not globally static) asymptotically AdS geometries for which the RT and HRT proposals might differ, depending on the precise formulation of the homology constraint and of RT’s constant time slice. We have not offered any definitive resolution; this is an interesting and important area which we leave for future exploration. The specific simple examples we have presented above all have some potentially dubious features: in case of Reissner-Nordstrom-AdS, the Cauchy horizon is unstable, while in the other two cases, the shell is unstable. Moreover, in all cases, the difference between the prescriptions occurs due to a part of the geometry which is beyond the horizon and thus causally inaccessible to an asymptotic observer. While this feature might therefore seem rather unappealing, we stress that in general we would be forced into such situation in any case, as long as the entanglement entropy is related to

²⁴ However, as indicated in §7, similar effect takes place in time-dependent situations involving black hole collapse.

some locally-defined geometric construct, due to the teleological nature of the event horizon. The examples mentioned above merely illustrate the issues we have yet to confront to fully understand the holographic entanglement entropy prescription.

7 Discussion

We have focused on exploring the behaviour of entanglement entropy $S_{\mathcal{A}}$ under smooth deformations of the entangling region \mathcal{A} in finite systems. Of particular interest to us is the distinction between the behaviour of holographic field theories, which in the large c (planar limit) and strong coupling limit can be mapped onto the dynamics of classical gravity, versus field theories away from the planar limit. Sharp features in observables are possible in the latter since the planar limit allows one to enter a ‘thermodynamic regime’, as evidenced for example by the thermal phase transitions in finite volume [37–39].

We focused in particular on the Araki-Lieb inequality which gives a useful measure of the relative entanglement of a region \mathcal{A} and its complement \mathcal{A}^c . While this would vanish if the entire system were in a pure state, it carries non-trivial information about the total density matrix in general. We indeed encounter an interesting phenomenon of entanglement plateaux: the Araki-Lieb inequality is saturated for finite system sizes, owing to some non-trivial features of minimal surfaces which compute the holographic entanglement entropy.

Focusing on thermal density matrices in CFTs we find that in $1 + 1$ dimensional field theory the Araki-Lieb inequality forces us to modify the expression of entanglement entropy for a large enough region (a point previously noted in [19, 20]) and provide an analytic prediction of when the plateau is attained. We also contrast this behaviour of large c theories against low central charge theories. There being very few exact results on the entanglement entropy of thermal CFTs in finite volume, we focused on the available expressions for Dirac fermion (with and without chemical potential). In the $c = 1$ case we note the absence of any plateau and the Araki-Lieb inequality is only saturated when the region under consideration (or its complement) is maximal. In higher dimensional holographic examples this no longer is the case, $\delta S_{\mathcal{A}}$ is forced by virtue of the features of the holographic construction to plateau (see below).

For the main part of this work, we have considered the RT prescription for calculating the holographic entanglement entropy, which is valid for static equilibrium situations. In this context, one is instructed to work at a given instant in time, and the entanglement entropy computation then involves finding the area of a requisite *minimal* surface in the corresponding Riemannian geometry. Though innocuously simple-sounding, there is a rich set of features associated with what precisely is meant by ‘requisite’. For a specified boundary region \mathcal{A} , the boundary of the relevant surface $\mathfrak{M}_{\mathcal{A}}$ must coincide with the entangling surface $\partial\mathcal{A}$, it must be homologous to \mathcal{A} , and in case of multiple such surfaces, it must be the one with smallest area.

It has already been observed previously that the last restriction can cause the entanglement entropy to have a kink (i.e., its first derivative with respect to the parameter α characterizing

the region \mathcal{A} can be discontinuous). This is because there can exist multiple families of minimal surfaces which can exchange dominance. Here we have explored this multiplicity further, and discovered that even in the most simple case of global eternal Schwarzschild-AdS $_{d+1}$ geometry with $d \geq 3$, in a certain regime of α there is actually an infinite tower of minimal surfaces anchored on the same entangling surface (though only the lowest two have a regime of dominance), seemingly approaching self-similar behaviour. This curious feature can arise thanks to the dimensionality of the surfaces and compactness of the horizon. On the other hand, in other regimes of α (namely for sufficiently large region \mathcal{A}), there is only a single, disconnected minimal surface satisfying the homology requirement – unlike in the $2 + 1$ dimensional case, a connected minimal surface homologous to \mathcal{A} simply does not exist. This novel feature may be understood as a consequence of certain properties of causal wedges discussed in [23].

We note in passing that the distinction between the AdS $_3$ and higher dimensional examples is quite reminiscent of the Hagedorn behaviour encountered in the dual field theories. In AdS $_3$ the BTZ black hole saddle point exists for all values of temperature (as does the thermal AdS one), but it becomes sub-dominant at low temperatures. In higher dimensions, the Schwarzschild-AdS saddles only exist above a minimum temperature T_{min} ; they however take over from the thermal AdS saddle at a slightly higher temperature $T_H > T_{min}$ (when the horizon size is comparable to AdS length scale). So at low enough temperatures we are always forced into the confining state. In the context of entanglement entropy the analogous observation is about the existence of connected minimal surfaces in the black hole geometry (we should emphasize that we are always in the deconfined phase to be able to use the black hole saddle). In three dimensions, connected surfaces always exist but fail to be dominant past some critical region size; in higher dimensions they cease to exist past a critical region size given the homology constraint.

Our discussion was primarily focussed on entanglement entropy and in particular on $\delta S_{\mathcal{A}}$. One could equally have focussed on the mutual information, $I(\mathcal{A}, \mathcal{B}) = S_{\mathcal{A}} + S_{\mathcal{B}} - S_{\mathcal{A} \cup \mathcal{B}}$, for two disjoint regions \mathcal{A} and \mathcal{B} . Holographic studies of mutual information also reveal an interesting behaviour: for sufficiently separated regions $I(\mathcal{A}; \mathcal{B}) = 0$. This is due to the fact that the extremal surface $\mathfrak{E}_{\mathcal{A} \cup \mathcal{B}}$ breaks up into a disconnected surface anchored on \mathcal{A} and \mathcal{B} to avoid thin necks. This has been extensively discussed in the literature cf., [6] and [31, 33, 40] for implications of this construct in the context of probing black hole interiors. For the regions we have considered \mathcal{A} and \mathcal{A}^c are not strictly disjoint, so we have to be a bit more careful; naively the mutual information involves the sum of the entropies of two regions (which we take to be \mathcal{A} and \mathcal{A}^c) while the Araki-Lieb inequality constrains their difference, giving $I(\mathcal{A}; \mathcal{A}^c) = 2S_{\mathcal{A}}$ which is UV divergent.²⁵ One can however constrain the mutual information between the smaller of \mathcal{A} or \mathcal{A}^c and the auxiliary system $\hat{\rho}_{\Sigma}$ which purifies the global density matrix ρ_{Σ} (thus $S_{\rho_{\Sigma}} = S_{\mathcal{A} \cup \mathcal{A}^c} = S_{\hat{\rho}_{\Sigma}}$) to vanish when

²⁵ We thank the referee for spotting an error in our original statement and for suggesting use of the auxiliary purifying system.

Araki-Lieb inequality is saturated. For example one has

$$I(\mathcal{A}; \hat{\rho}_\Sigma) = S_{\mathcal{A}} + S_{\hat{\rho}_\Sigma} - S_{\mathcal{A} \cup \hat{\rho}_\Sigma} = S_{\mathcal{A}} + S_{\rho_\Sigma} - S_{\mathcal{A}^c} = \begin{cases} 0, & \text{for } \alpha < \frac{1}{2} \\ 2S_{\rho_\Sigma}, & \text{for } \alpha > \frac{1}{2} \end{cases} \quad (7.1)$$

In deriving this we used the fact that $S_{\mathcal{A} \cup \mathcal{A}^c \cup \hat{\rho}_\Sigma} = 0$ by definition of the purifying degrees of freedom. In the case of the global thermal density matrix $\hat{\rho}_\Sigma$ is the thermofield double which lives on the second asymptotic region of the eternal black hole Penrose diagram.²⁶

While these causal wedge considerations already by themselves guarantee that the entanglement entropy $S_{\mathcal{A}}(\alpha)$ cannot be a smooth function for thermal states in higher dimensions, we have argued that, if defined by a minimization procedure as in the RT prescription, the entanglement entropy must nevertheless be continuous. Our argument crucially used minimality. In particular, in the HRT prescription of computing entanglement entropy via a smallest area *extremal* surface, the continuity argument given in §2 does not apply. It would be interesting to explore whether sufficient time-dependence can provide counter-examples to continuity or whether the one can generalize the proof of continuity to the Lorentzian context; we leave further investigation of this issue for the future [12].

Note that in this collapsed black hole context, the homology requirement is satisfied automatically, though as explained in [12], that itself leads to very curious feature of entanglement entropy: it can distinguish between an eternal black hole and a collapsed one, arbitrarily long after the collapse had taken place [41]. In this sense, while in the AdS/CFT context we are used to classical bulk surfaces providing at best only some coarse-grained CFT information, the homology requirement induces a fine-grained aspect to the entanglement entropy observable.

While in the time-dependent setting mentioned above, the RT prescription is meaningless and therefore cannot be compared to the extremal surface generalization of HRT, we have seen that there are subtle differences between the RT and HRT proposals even in static situations. Indeed, a similar ‘fine-grained’ quality as mentioned above is manifested in the entanglement entropy for e.g. the Schwarzschild-AdS bag of gold geometry discussed in §6: there the entanglement entropy (as given by the HRT prescription) can easily distinguish the differences in the geometry behind the horizon, on the other side of the Einstein-Rosen bridge. Other known diagnostics of the CFT density matrix not being precisely thermal typically require rather sophisticated machinery, such as detecting lack of periodicity on the Euclidean time circle of the analytically continued solution. We find it remarkable that entanglement entropy does the job so easily. This observation of course crucially hinges on the precise formulation of the homology constraint, which deserves to be understood better.

²⁶ It is worth mentioning that the saturation of the Araki-Lieb inequality in large c theories implies that the degrees of freedom in \mathcal{A} admit a canonical split into two uncorrelated parts: one that carries all the entanglement with \mathcal{A}^c and another that carries the macroscopic entropy of ρ_Σ , cf., [42]. We thank Matt Headrick for suggesting this interpretation and other comments regarding Araki-Lieb inequality saturation.

Acknowledgments

It is a pleasure to thank Jan de Boer, Juan Jottar, Hong Liu, Don Marolf, Tadashi Takayanagi and especially Matt Headrick for various illuminating discussions. We would also like to thank Hong Liu, Matthew Headrick and Tadashi Takayanagi for their detailed comments on a draft of this manuscript. VH and MR would like to thank CERN, ITF, Amsterdam and ICTP for hospitality during the early stages of this project. In addition MR would like to acknowledge the hospitality of Technion, Israel during the “Relativistic fluid dynamics and the gauge gravity duality” workshop as well as the organizers of the “2nd Mediterranean Conference on Classical And Quantum Gravity” Veli Losinj, Croatia for their hospitality during the concluding stages of this project. HM is supported by a STFC studentship. VH and MR are supported in part by the STFC Consolidated Grant ST/J000426/1.

A Mean Curvature Flow

In this appendix we summarize the algorithm we used to construct the minimal surfaces in Schwarzschild-AdS $_{d+1}$ for $d \geq 3$. While one can directly numerically solve the Euler-Lagrange equations of (4.2) directly, it is especially useful to consider alternate strategies to ensure that we have obtained all the surfaces of interest.

For conceptual and computational purposes, it is useful to view minimal surfaces as the endpoint of a flow of surfaces of decreasing area. This is analogous to a gradient descent algorithm for finding a local minimum of a function, for which one considers a point moving with velocity equal to minus the gradient of the function. This gives a curve along which the function monotonically decreases, and exponentially decays to the minimum as a function of flow time $t_{flow} \equiv \mathfrak{t} \rightarrow \infty$.

In the analogous process for minimal submanifolds, the function to be minimized becomes the area functional, and the flow velocity vector becomes a vector field (which may be taken to be normal) defined on the submanifold. The appropriate vector field is given by the *mean curvature* of the submanifold, and the evolutionary process is known as mean curvature flow.

We review the required differential geometric technology, and outline an algorithm for computing the process in the case of codimension-one minimal surfaces.²⁷

A.1 Some geometry

Let \mathfrak{n} be a smooth n -dimensional Riemannian manifold (e.g., \mathfrak{n} for situations of interest is a constant time surface in the bulk spacetime \mathcal{M}), with metric g and Levi-Civita connection ∇ . Further, let $\mathfrak{m} \subset \mathfrak{n}$ be a smooth m -dimensional submanifold (with $m < n$), defined by embedding

²⁷ The minimal surfaces relevant for computing entanglement entropy using the prescription of [1] are co-dimension two in the bulk. In static spacetimes however we can localize on a constant time bulk hypersurface, leading thus to a search of co-dimension one surfaces on the slice.

$\psi : \mathfrak{m} \rightarrow \mathfrak{n}$. The induced metric on \mathfrak{m} is inherited from \mathfrak{n} as $\gamma = \psi^*g$ given this embedding. Vectors at a point on the submanifold can be decomposed into components tangent and normal to \mathfrak{m} : $U = U^\top + U^\perp$.

The natural structures associated with $\mathfrak{m} \subset \mathfrak{n}$ are in terms of the intrinsic and extrinsic geometry of the embedding, characterized by the induced metric γ (the *first fundamental form*) and the extrinsic curvature K (the *second fundamental form*). Formally, given vector fields, say U and V , tangent to \mathfrak{m} , one can decompose the covariant derivative:

$$\nabla_U V = (\nabla_U V)^\top + (\nabla_U V)^\perp = \mathcal{D}_U V + K(U, V), \quad (\text{A.1})$$

which effectively defines \mathcal{D} and K . The former is the Levi-Civita connection of the induced metric on \mathfrak{m} , while the latter is a normal-valued symmetric bi-linear form. It is worth emphasizing that the second fundamental form depends only on the vectors at the point, so it is a tensor on \mathfrak{m} . A particularly important object is obtained by taking its trace \mathbf{k} , giving a normal vector field on \mathfrak{m} : this is the *mean curvature field*.²⁸ In components, using coordinates x^a intrinsic to \mathfrak{m} , and X^α on \mathfrak{n} , we have $\gamma_{ab} = g_{\alpha\beta} \partial_a X^\alpha \partial_b X^\beta$, and

$$h^{\alpha\beta} = \gamma^{ab} \partial_a X^\alpha \partial_b X^\beta, \quad K_{\alpha\beta}{}^\gamma = h_\alpha{}^\mu h_\beta{}^\nu \nabla_\nu h_\mu{}^\gamma, \quad \mathbf{k}^\alpha = h^{\mu\nu} K_{\mu\nu}{}^\alpha.$$

Here h is the projection onto vectors tangent to \mathfrak{m} . Note that the extrinsic curvature tensor $K_{\alpha\beta}{}^\gamma$ is symmetric in its lower indices which lie tangent to \mathfrak{m} , while the upper index is transverse.

We are now in a position to make precise the geometric interpretation of the mean curvature field alluded to above, that it is minus the gradient of the area functional on the space of submanifolds.

Let us consider a smooth family of embeddings $\Psi : I \times \mathfrak{m} \rightarrow \mathfrak{n}$, where I is some interval parameterized by a flow time \mathfrak{t} . For each $\mathfrak{t} \in I$, $\psi_\mathfrak{t} = \Psi(\mathfrak{t}, \cdot) : \mathfrak{m} \rightarrow \mathfrak{n}$ is an embedding. Let A be the area of the embedded submanifold:

$$A(\mathfrak{t}) = \int_{\mathfrak{m}} \omega_\mathfrak{t}, \quad (\text{A.2})$$

where $\omega_\mathfrak{t}$ is the induced volume form on \mathfrak{m} at time \mathfrak{t} .

Let V be the vector field given by the local velocity of the surface in this deformation, so $V = \Psi_* \frac{\partial}{\partial \mathfrak{t}}$. One should think of Ψ , or $\psi_\mathfrak{t}$, as defining a curve through the space of embeddings, an infinite-dimensional manifold whose tangent space at ψ is given by the set of $T\mathfrak{n}$ -valued vector fields on ψ . Then V is naturally the tangent vector to the curve.

We now have enough state our key result:

$$\frac{dA}{d\mathfrak{t}} = - \int_{\mathfrak{m}} g(\mathbf{k}, V) \omega_\mathfrak{t}. \quad (\text{A.3})$$

²⁸ This definition of the mean curvature is not entirely universal: in some conventions the sign differs, while in others the trace is divided by a factor of m , so that \mathbf{k} is genuinely the mean of the eigenvalues of K (which incidentally define the principal curvatures of the submanifold).

In the space of embeddings, this is the natural inner product between V and $-\mathbf{k}$, which makes precise the analogy with a gradient. Note that, as is to be expected, only the normal component of V contributes, since the tangent component is ‘pure gauge’, describing how parameterization of the embedding evolves, and not the shape of the submanifold itself.

Proof. The method of proof is to compute $\frac{\partial}{\partial \mathbf{t}}$ of the induced volume form on \mathbf{m} , and we will show it is equal to $(-g(\mathbf{k}, V) + \mathcal{D} \cdot V^\top) \omega_{\mathbf{t}}$. Since the last term is a total derivative, the result will follow, provided that the variation is constrained to a compact region, or we restrict the variation to be normal (so $V^\top = 0$). For ease of computation, we work in a chart of normal coordinates x^i on \mathbf{m} at the point p , with respect to the induced metric at time $\mathbf{t} = 0$: $\gamma(\partial_i, \partial_j) = \delta_{ij}$, and $\mathcal{D}_{\partial_i} \partial_j = 0$. Let $e_i = \Psi_* \partial_i$. For later use, we note that since Lie brackets commute with push-forward, $[e_i, V] = 0$, so $\nabla_V e_i = \nabla_{e_i} V$.

The induced volume form is given by $\omega_{\mathbf{t}} = \sqrt{\det(g(e_i, e_j))} dx^1 \wedge \cdots \wedge dx^m$, so its derivative at the point p and $\mathbf{t} = 0$ is

$$\frac{1}{2} \frac{\partial}{\partial \mathbf{t}} (g(e_i, e_i)) \omega_0, \quad (\text{A.4})$$

where we have used the fact that the derivative of $\det A$ is $\det A \operatorname{tr}(A^{-1} \dot{A})$, and that $dx^1 \wedge \cdots \wedge dx^m = \omega_0$ at p in our normal coordinates. Now, again computing at the point p and $\mathbf{t} = 0$,

$$\begin{aligned} \frac{1}{2} \frac{\partial}{\partial \mathbf{t}} g(e_i, e_i) &= \frac{1}{2} V(g(e_i, e_i)) \\ &= g(\nabla_V e_i, e_i) \\ &= g(\nabla_{e_i} V, e_i) \\ &= e_i(g(V, e_i)) - g(V, \nabla_{e_i} e_i) \\ &= \partial_i V_i^\top - g(V, \mathcal{D}_{e_i} e_i + K(e_i, e_i)) \\ &= \mathcal{D} \cdot V^\top - g(V, \mathbf{k}). \end{aligned}$$

In the final line we have used the fact that we have normal coordinates, so $\partial_i = \mathcal{D}_i$, and the sum over i reduces to the trace. \square

A.2 The algorithm

With the necessary mathematical formalism in hand, we now describe an algorithm for implementing the mean curvature flow, specializing for simplicity to the case of codimension-one surfaces. In this case, having chosen a direction for a normal, the mean curvature vector field reduces to a scalar field, since the normal space is one-dimensional.

The method used is based on a level-set approach: the surfaces are described by the sets

$$\mathbf{m} = \{x \in \mathbf{n} \mid \Phi(x) = c\} \quad (\text{A.5})$$

for some function Φ and some constant c , and the evolution of the function Φ is what we shall model. This has several advantages:

- There is no requirement to pick a parameterization of the surface, so no problem with ‘gauge fixing’.
- A whole family of surfaces can be modeled at once, by computing the evolution for Φ and picking different level sets.
- The algorithm elegantly allows for changes in topology: under mean curvature flow, surfaces can split into several parts, and this causes no trouble for the level set method.

The definition of the flow gives the normal speed of a point on the surface as it evolves: $n(\frac{dx}{dt}) = \mathbf{k}$, with n being the normal to \mathbf{m} . Provided that $d\Phi \neq 0$, we also have a simple way of computing a unit normal, $n = \frac{d\Phi}{\|d\Phi\|}$. To deduce the evolution of Φ , differentiate the definition (A.5) of a surface \mathbf{m} :

$$0 = \frac{\partial \Phi}{\partial \mathbf{t}} + d\Phi\left(\frac{dx}{dt}\right) = \frac{\partial \Phi}{\partial \mathbf{t}} + \|d\Phi\| n\left(\frac{dx}{dt}\right) = \frac{\partial \Phi}{\partial \mathbf{t}} + \|d\Phi\| \mathbf{k}.$$

The only ingredient which remains is to compute \mathbf{k} using only information from Φ . Luckily, this is given by a simple formula in terms of the normal $\mathbf{k} = -\nabla \cdot n$. To see this, take n to be the unit one-form, normal to the surface, but also defined, and unit, off it. Extend to an orthonormal basis²⁹ $\{e_i, n^\#\}$. Then:

$$\begin{aligned} \mathbf{k} &= n(\nabla_{e_i} e_i) = e_i(n(e_i)) - (\nabla_{e_i} n)(e_i) \\ &= -\nabla \cdot n + (\nabla_{n^\#} n)(n^\#) \\ &= -\nabla \cdot n + g(n, \nabla_{n^\#} n) \\ &= -\nabla \cdot n + \frac{1}{2} n^\#(g(n, n)) = -\nabla \cdot n. \end{aligned}$$

The term $e_i(n(e_i))$ in the first line vanishes because $n(e_i) = 0$ everywhere, and the last equality follows because $g(n, n) = 1$ everywhere.

This gives us our final evolution equation for Φ :

$$\frac{\partial \Phi}{\partial \mathbf{t}} = \|d\Phi\| \nabla \cdot \left(\frac{d\Phi}{\|d\Phi\|} \right) = \sqrt{g^{\gamma\delta} \partial_\gamma \Phi \partial_\delta \Phi} \frac{1}{\sqrt{g}} \partial_\alpha \left(\frac{\sqrt{g} g^{\alpha\beta} \partial_\beta \Phi}{\sqrt{g^{\gamma\delta} \partial_\gamma \Phi \partial_\delta \Phi}} \right). \quad (\text{A.6})$$

This is a quasilinear parabolic PDE in Φ , which can be solved numerically by standard methods.

For the computations in this paper, only spherically symmetric surfaces were considered, so Φ was a function of two spatial coordinates θ and r (or another suitably redefined radial coordinate). The PDE was solved with difference methods, with Neumann boundary conditions at the poles of the sphere ($\theta = 0, \pi$) to keep the surfaces regular there, and mixed boundary conditions at radial boundaries set by the desired asymptotics.

²⁹ We denote the ‘raised index’ normal vector by $n^\#$, so that $g(n^\#, U) = n(U)$

B Near-horizon behaviour of minimal surfaces

An analytic solution of the equations of motion in a near-horizon approximation explains the existence of an infinite self-similar family of minimal surfaces, wrapping the horizon multiple times. We find that this is generic near any minimum of the radius of the \mathbf{S}^{d-1} , such as the Einstein-Rosen bridge of a non-extremal black hole, and is characterized entirely by the number of dimensions,³⁰ and a single parameter $\xi \equiv r_+ \kappa$, where κ is the surface gravity (which is related to the black hole temperature as $T = \frac{1}{4\pi} \kappa$).

For the analysis we use a modified radial coordinate z , regular at the event horizon, defined by $r = r_+ + \frac{1}{2} r_+ \xi z^2$, and parametrize the surface by θ throughout. Dots will denote differentiation with respect to θ .

We invoke two approximation schemes. The naive linearization holds away from the poles, but breaks down close the $\theta = 0, \pi$ where \dot{z} grows without bound. Near the poles, we use a different approximation, reducing the problem locally to flat Euclidean space. We will find that these two approximation schemes have overlapping domains of validity, so the complete behaviour can be described by matching the asymptotics in each regime. This predicts a tower of arbitrarily many branches of the surface wrapping the horizon, with a known relationship between the distances from the horizon of successive surfaces.

Schematically: we start at the North pole $\theta = 0$, and solve in the linear regime towards the South pole $\theta = \pi$. There we change to the near-pole regime, in which there will be a turning point at a minimum angle Δ away from the pole, at $z = z_0$. These constants are both computed from matching to the linearized solution. We then match up to another branch moving back towards the North pole, where there will be a turning point given by new, parametrically larger, Δ and z_0 . The matching process will allow us to produce a recurrence relation for the sequences Δ_n, z_n of turning points. We can obtain an entirely analytic expression for these values, which, sufficiently close to the horizon, match precisely with the values obtained from numerically integrated surfaces.

B.1 The linear regime

We start with the linear approximation, valid for $z, \dot{z} \ll 1$. In this case, the minimal surface equation obtained from (4.2) reduces to

$$\ddot{z} + (d-2) \cot \theta \dot{z} - (d-1) \xi z = 0, \quad \xi \equiv r_+ \kappa. \quad (\text{B.1})$$

This equation has regular singular points at $\theta = 0, \pi$, where the solutions have asymptotics $z \sim \theta^\sigma$ for $\sigma = 0, 3-d$ (or constant and $z \sim \log \theta$ asymptotics in $d = 3$). For $\xi > 0$ it is easy to show that a solution regular at $\theta = 0$ cannot be regular at $\theta = \pi$, so the general solution can be written

$$z(\theta) = A z_r(\theta) + B z_r(\pi - \theta) \quad (\text{B.2})$$

³⁰ Here we will keep to spatial codimension-one surfaces, but the argument carries over essentially unchanged for surfaces of any dimension m , wrapping a \mathbf{S}^{m-1} in an $SO(d)$ symmetric space

using the symmetry under $\theta \mapsto \pi - \theta$. Here z_r is chosen by demanding $z_r(\theta) \rightarrow 1$ as $\theta \rightarrow 0$, and the asymptotics are characterized by a parameter λ , by $z_r(\theta) \sim \lambda (\pi - \theta)^{3-d}$ as $\theta \rightarrow \pi$. In fact, the solutions can be found explicitly in terms of associated Legendre functions, and $\lambda(d, \xi)$ can be found analytically.

B.2 Near-pole flat space regime

For the second approximation scheme, we assume $z \ll 1, z \ll \dot{z}, \theta \ll 1$, in which case the problem reduces to that of a spherically symmetric “soap bubble” in flat d -dimensional Euclidean space, where (θ, z) are the radial and longitudinal coordinates respectively in a cylindrical polar system. This has the solution

$$z = z_0 \pm \Delta \int_1^{\frac{\theta}{\Delta}} \frac{dx}{\sqrt{x^{2d-4} - 1}}. \quad (\text{B.3})$$

The constants z_0 and Δ will be small values, describing the distance from horizon and pole respectively of the turning point of the surface. Note that taking the limit $\theta \rightarrow \infty$, the integral converges, so z tends to a fixed value when $d > 3$: the upper and lower branches of the surface remain at bounded separation, characterized by the value of

$$I_d = \int_1^\infty \frac{dx}{\sqrt{x^{2d-4} - 1}} = -\sqrt{\pi} \frac{\Gamma\left(\frac{d-3}{2(d-2)}\right)}{\Gamma\left(\frac{-1}{2(d-2)}\right)}. \quad (\text{B.4})$$

The asymptotics are given by

$$z \sim z_0 \pm I_d \Delta \mp \frac{1}{d-3} \Delta^{d-2} \theta^{3-d} \quad \text{as} \quad \frac{\theta}{\Delta} \rightarrow \infty. \quad (\text{B.5})$$

B.3 Matching asymptotics

With solutions in the two approximation schemes in hand, it now remains only to fix constants of integration by matching the asymptotics. For definiteness, consider starting from a known (Δ_n, z_n) close to the North pole, computing the coefficients of the linear solution, and finally extracting the values of (Δ_{n+1}, z_{n+1}) close to the South pole.

The case $d = 3$ is a little different from $d > 3$, because of the logarithmic asymptotics, so we will stick to $d > 3$ here. A very similar analysis holds in this case, for which we will omit the details, and present only the results.

Recall the form of the linear solution $z(\theta) = A z_r(\theta) + B z_r(\pi - \theta)$, with the terms regular at $\theta = 0, \pi$ respectively. Consistency imposes that the coefficients have parametrically different sizes, with $A \gg B$. Towards the South pole, the B -mode is also suppressed relative to the A -mode by the growth of the latter, which renders B essentially irrelevant, and the turning points can be characterized by a single parameter. Indeed, by matching the asymptotics near the South pole, for consistency we need $z_{n+1} = I_d \Delta_{n+1}$. Finally, we match A with the relevant coefficients at each

end,

$$z_n + I_d \Delta_n = A = \frac{\Delta_{n+1}^{d-2}}{\lambda(d-3)}, \quad (\text{B.6})$$

which finally gives us our recurrence relation

$$\Delta_{n+1} = [2\lambda(d-3)\Delta_n]^{\frac{1}{d-2}}, \quad (\text{B.7})$$

demonstrating a parametric growth in Δ, z at each step. This can be conveniently solved by changing to the coordinate x , defined by $r = r_+ + e^{2x}$, (or $\frac{1}{2}r_+ \xi z^2 = e^{2x}$) from which we get

$$x_n = \frac{x_0 - \nu}{(d-2)^n} + \nu, \quad \Delta_n = \sqrt{\frac{2}{\kappa}} \frac{e^{x_n}}{r_+ I_d} \quad (\text{B.8})$$

with the constant ν given by

$$\nu = \frac{1}{2} \log\left(\frac{r_+ \xi}{2}\right) + \frac{\log[2(d-2)\lambda I_d^{d-2}]}{d-3}. \quad (\text{B.9})$$

It should be noted that the initial value of x where the surface smoothly crosses the pole $\theta = 0$, closest to the horizon, is not at x_0 , but rather at $x = x_0 + \log 2$, since the solution here is not given by (B.3), but by that obtained in the linear regime: constant x to leading order.

For completeness, we record the computed values of λ here:

$$\begin{aligned} \lambda &= \frac{\cosh(\pi j)}{(d-3)\pi} \frac{2^{d-2} \left(\left(\frac{d-3}{2}\right)!\right)^2}{\left[\left(\frac{1}{2}\right)^2 + j^2\right] \left[\left(\frac{3}{2}\right)^2 + j^2\right] \cdots \left[\left(\frac{d-4}{2}\right)^2 + j^2\right]}, & d \text{ odd} \\ \lambda &= (d-3) \frac{\sinh(\pi j)}{j} \left(\frac{1}{1^2 + j^2}\right) \left(\frac{3}{2^2 + j^2}\right) \left(\frac{5}{3^2 + j^2}\right) \cdots \left(\frac{d-5}{\left(\frac{d-4}{2}\right)^2 + j^2}\right), & d \text{ even} \end{aligned} \quad (\text{B.10})$$

where in both cases, j is defined by

$$j^2 = (d-1)\xi - \frac{(d-2)^2}{4}. \quad (\text{B.11})$$

In $d = 3$, we have slightly altered formulae:

$$\Delta_{n+1} = -2\lambda\Delta_n \log\left(\frac{\Delta_n}{2}\right) \quad (\text{B.12})$$

$$x_n = \frac{1}{2} \log\left(\frac{r_+ \xi}{2}\right) + \log\left[-\Delta_n \log\left(\frac{\Delta_n}{2}\right)\right] \quad (\text{B.13})$$

$$\lambda = \frac{2 \cosh(\pi j)}{\pi}. \quad (\text{B.14})$$

B.4 Validity

Having done this calculation, it is worth checking to see that we have done something sensible: there must be some overlap between the domains of validity of each approximation. It is straightforward to confirm this *a posteriori*, looking at overlaps near the North pole for definiteness.

Firstly, the linear approximation is valid when $z, \dot{z} \ll 1$, which requires θ^{2-d} to be much smaller than the coefficient of the growing mode. This holds for $\theta \gg \Delta$. Secondly, the near-pole approximation is valid when $z \ll \dot{z}$, which requires $\theta^{d-2} \ll \Delta^{d-3}$ (as well as $\theta \ll 1$). This means that as long as Δ remain small, both approximations are simultaneously valid for

$$\Delta \ll \theta \ll \Delta^{\frac{d-3}{d-2}}. \quad (\text{B.15})$$

For $d = 3$, we have a similar result, with overlap region

$$e^{-\frac{1}{\Delta}} \ll \theta \ll \frac{-1}{\log \Delta}. \quad (\text{B.16})$$

B.5 Relation to θ_∞

We can straightforwardly use this analysis to show that for θ_∞ sufficiently close to $\pi/2$ there will be an infinite tower of minimal surfaces. This is because each branch of a surface is characterized by a single parameter. All surfaces with a particular turning point will match closely outside that turning point. To obtain a surface on later branches of this tower requires the initial value of $-x$ to be exponentially large, so the distance from the horizon will be suppressed by an exponential of an exponential.

Incidentally, this gives an indirect demonstration of non-existence of connected minimal surfaces for sufficiently large θ_∞ . The connected surfaces form a one-parameter family, parameterized by the maximum depth reached into the bulk, or equivalently the location x_0 of the surface at $\theta = 0$. Considering θ_∞ as a function of x_0 , continuity implies that there are only two ways for θ_∞ to reach arbitrarily close to π . The first is for $\theta_\infty = \pi$ for some x_0 , but this is precluded from happening by behaviour near the boundary. The other way is for a sequence x_m to exist such that $\theta_\infty(x_m)$ tends to π , but the self-similarity properties of surfaces indicated here will prevent this option also.

References

- [1] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Phys. Rev. Lett.* **96** (2006) 181602, [[hep-th/0603001](#)].
- [2] S. Ryu and T. Takayanagi, *Aspects of holographic entanglement entropy*, *JHEP* **08** (2006) 045, [[hep-th/0605073](#)].
- [3] V. E. Hubeny, M. Rangamani, and T. Takayanagi, *A Covariant holographic entanglement entropy proposal*, *JHEP* **0707** (2007) 062, [[arXiv:0705.0016](#)].
- [4] R. Bousso, *The Holographic principle*, *Rev.Mod.Phys.* **74** (2002) 825–874, [[hep-th/0203101](#)].
- [5] D. V. Fursaev, *Proof of the holographic formula for entanglement entropy*, *JHEP* **0609** (2006) 018, [[hep-th/0606184](#)].
- [6] M. Headrick, *Entanglement Renyi entropies in holographic theories*, *Phys.Rev.* **D82** (2010) 126010, [[arXiv:1006.0047](#)].

- [7] H. Casini, M. Huerta, and R. C. Myers, *Towards a derivation of holographic entanglement entropy*, *JHEP* **1105** (2011) 036, [[arXiv:1102.0440](#)].
- [8] A. Lewkowycz and J. Maldacena, *Generalized gravitational entropy*, [arXiv:1304.4926](#).
- [9] T. Hartman, *Entanglement Entropy at Large Central Charge*, [arXiv:1303.6955](#).
- [10] T. Faulkner, *The Entanglement Renyi Entropies of Disjoint Intervals in AdS/CFT*, [arXiv:1303.7221](#).
- [11] H. Araki and E. Lieb, *Entropy inequalities*, *Commun.Math.Phys.* **18** (1970) 160–170.
- [12] V. E. Hubeny and H. Maxfield, *work in progress*.
- [13] T. Albash and C. V. Johnson, *Holographic Studies of Entanglement Entropy in Superconductors*, *JHEP* **1205** (2012) 079, [[arXiv:1202.2605](#)].
- [14] A. Belin, A. Maloney, and S. Matsuura, *Holographic Phases of Renyi Entropies*, [arXiv:1306.2640](#).
- [15] T. Nishioka and T. Takayanagi, *AdS Bubbles, Entropy and Closed String Tachyons*, *JHEP* **0701** (2007) 090, [[hep-th/0611035](#)].
- [16] I. R. Klebanov, D. Kutasov, and A. Murugan, *Entanglement as a probe of confinement*, *Nucl.Phys.* **B796** (2008) 274–293, [[arXiv:0709.2140](#)].
- [17] H. Liu and M. Mezei, *A Refinement of entanglement entropy and the number of degrees of freedom*, [arXiv:1202.2070](#).
- [18] M. Headrick and T. Takayanagi, *A Holographic proof of the strong subadditivity of entanglement entropy*, *Phys.Rev.* **D76** (2007) 106013, [[arXiv:0704.3719](#)].
- [19] T. Azeyanagi, T. Nishioka, and T. Takayanagi, *Near Extremal Black Hole Entropy as Entanglement Entropy via AdS(2)/CFT(1)*, *Phys.Rev.* **D77** (2008) 064005, [[arXiv:0710.2956](#)].
- [20] D. D. Blanco, H. Casini, L.-Y. Hung, and R. C. Myers, *Relative Entropy and Holography*, [arXiv:1305.3182](#).
- [21] B. Freivogel, V. E. Hubeny, A. Maloney, R. C. Myers, M. Rangamani, *et. al.*, *Inflation in AdS/CFT*, *JHEP* **0603** (2006) 007, [[hep-th/0510046](#)].
- [22] V. E. Hubeny and M. Rangamani, *Causal Holographic Information*, *JHEP* **1206** (2012) 114, [[arXiv:1204.1698](#)].
- [23] V. E. Hubeny, M. Rangamani, and E. Tonni, *Global properties of causal wedges in asymptotically AdS spacetimes*, [arXiv:1306.4324](#).
- [24] P. Calabrese and J. L. Cardy, *Entanglement entropy and quantum field theory*, *J.Stat.Mech.* **0406** (2004) P06002, [[hep-th/0405152](#)].
- [25] N. Ogawa, T. Takayanagi, and T. Ugajin, *Holographic Fermi Surfaces and Entanglement Entropy*, *JHEP* **1201** (2012) 125, [[arXiv:1111.1023](#)].
- [26] C. P. Herzog and M. Spillane, *Tracing Through Scalar Entanglement*, *Phys.Rev.* **D87** (2013) 025012, [[arXiv:1209.6368](#)].
- [27] M. M. Caldarelli, O. J. Dias, R. Emparan, and D. Klemm, *Black Holes as Lumps of Fluid*, *JHEP* **0904** (2009) 024, [[arXiv:0811.2381](#)].

- [28] R. Emparan, R. Suzuki, and K. Tanabe, *The large D limit of General Relativity*, [arXiv:1302.6382](#).
- [29] V. E. Hubeny, *Extremal surfaces as bulk probes in AdS/CFT*, *JHEP* **1207** (2012) 093, [[arXiv:1203.1044](#)].
- [30] J. Abajo-Arriastia, J. Aparicio, and E. Lopez, *Holographic Evolution of Entanglement Entropy*, *JHEP* **1011** (2010) 149, [[arXiv:1006.4090](#)].
- [31] T. Hartman and J. Maldacena, *Time Evolution of Entanglement Entropy from Black Hole Interiors*, *JHEP* **1305** (2013) 014, [[arXiv:1303.1080](#)].
- [32] H. Liu and S. J. Suh, *Entanglement Tsunami: Universal Scaling in Holographic Thermalization*, [arXiv:1305.7244](#).
- [33] S. H. Shenker and D. Stanford, *Black holes and the butterfly effect*, [arXiv:1306.0622](#).
- [34] A. C. Wall, *Maximin Surfaces, and the Strong Subadditivity of the Covariant Holographic Entanglement Entropy*, [arXiv:1211.3494](#).
- [35] B. Czech, J. L. Karczmarek, F. Nogueira, and M. Van Raamsdonk, *The Gravity Dual of a Density Matrix*, *Class.Quant.Grav.* **29** (2012) 155009, [[arXiv:1204.1330](#)].
- [36] D. Marolf, *Black Holes, AdS, and CFTs*, *Gen.Rel.Grav.* **41** (2009) 903–917, [[arXiv:0810.4886](#)].
- [37] E. Witten, *Anti-de Sitter space, thermal phase transition, and confinement in gauge theories*, *Adv.Theor.Math.Phys.* **2** (1998) 505–532, [[hep-th/9803131](#)].
- [38] B. Sundborg, *The Hagedorn transition, deconfinement and $N=4$ SYM theory*, *Nucl.Phys.* **B573** (2000) 349–363, [[hep-th/9908001](#)].
- [39] O. Aharony, J. Marsano, S. Minwalla, K. Papadodimas, and M. Van Raamsdonk, *The Hagedorn - deconfinement phase transition in weakly coupled large N gauge theories*, *Adv.Theor.Math.Phys.* **8** (2004) 603–696, [[hep-th/0310285](#)].
- [40] I. A. Morrison and M. M. Roberts, *Mutual information between thermo-field doubles and disconnected holographic boundaries*, [arXiv:1211.2887](#).
- [41] T. Takayanagi and T. Ugajin, *Measuring Black Hole Formations by Entanglement Entropy via Coarse-Graining*, *JHEP* **1011** (2010) 054, [[arXiv:1008.3439](#)].
- [42] L. Zhang and J. Wu, *On Conjectures of Classical and Quantum Correlations in Bipartite States*, *J.Phys.A: Math. Theor.* **45** (2012) 025301. [[arXiv:1105.2993](#)].