

UC Davis

UC Davis Previously Published Works

Title

Causality & holographic entanglement entropy

Permalink

<https://escholarship.org/uc/item/8pn2d5c9>

Journal

Journal of High Energy Physics, 2014(12)

ISSN

1126-6708

Authors

Headrick, Matthew
Hubeny, Veronika E
Lawrence, Albion
[et al.](#)

Publication Date

2014-12-01

DOI

10.1007/jhep12(2014)162

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-ShareAlike License, available at <https://creativecommons.org/licenses/by-nc-sa/4.0/>

Peer reviewed

Causality & holographic entanglement entropy

**Matthew Headrick^a, Veronika E. Hubeny^b, Albion Lawrence^a,
Mukund Rangamani^b**

^a*Martin Fisher School of Physics, Brandeis University,
MS 057, 415 South Street, Waltham, MA 02454, USA.*

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

E-mail: headrick@brandeis.edu, veronika.hubeny@durham.ac.uk,
albion@brandeis.edu, mukund.rangamani@durham.ac.uk

ABSTRACT: We identify conditions for the entanglement entropy as a function of spatial region to be compatible with causality in an arbitrary relativistic quantum field theory. We then prove that the covariant holographic entanglement entropy prescription (which relates entanglement entropy of a given spatial region on the boundary to the area of a certain extremal surface in the bulk) obeys these conditions, as long as the bulk obeys the null energy condition. While necessary for the validity of the prescription, this consistency requirement is quite nontrivial from the bulk standpoint, and therefore provides important additional evidence for the prescription. In the process, we introduce a codimension-zero bulk region, named the entanglement wedge, naturally associated with the given boundary spatial region. We propose that the entanglement wedge is the most natural bulk region corresponding to the boundary reduced density matrix.

KEYWORDS: AdS-CFT correspondence, Entanglement entropy

Contents

1	Introduction	1
2	Causal domains and entanglement entropy	6
2.1	Causality of entanglement entropy in QFT	6
2.2	Bulk geometry and holographic entanglement entropy	8
2.3	Causality constraints on extremal surfaces	12
2.4	Entanglement for disconnected boundary regions	14
3	Null geodesic congruences in AdS_3	17
3.1	Explicit solutions for geodesic congruences	18
3.2	Intersections within congruences	19
3.3	Expansion of congruences and caustics	21
3.4	Summary	23
4	Theorem and proof	24
4.1	Holographic setup	25
4.2	Congruences of null geodesics	26
4.3	Spatial regions and causal decompositions	29
5	Discussion	32

1 Introduction

One of the remarkable features of the holographic AdS/CFT correspondence is the geometrization of quantum-field-theoretic concepts. While certain aspects of recasting field-theory quantities into geometric notions have been ingrained in our thought, we are yet to fully come to grips with new associations between QFT and bulk geometry. A case in point is the fascinating connection of quantum entanglement and spacetime geometry. The genesis of this intricate and potentially deep connection harks back to the observation of Ryu-Takayanagi (RT) [1, 2] and subsequent covariant generalization by Hubeny-Rangamani-Takayanagi (HRT) [3] that the entanglement entropy of a quantum field theory is holographically computed by the area of a particular extremal surface in the bulk. In recent years, much effort has been expended in trying to

flesh out the physical implications of these constructions and in promoting the geometry/entanglement connection to a deeper level [4–7] which can be summarized rather succinctly in terms of the simple phrases “entanglement builds bridges” and “ER = EPR”. Whilst any connection between entanglement and geometry is indeed remarkable, further progress is contingent on the accuracy and robustness of this entry in the holographic dictionary. Let us therefore take stock of the status quo.¹

The RT proposal is valid for static states of a holographic field theory, which allows one to restrict attention to a single time slice $\tilde{\Sigma}$ in the bulk spacetime \mathcal{M} . The entanglement entropy of a region \mathcal{A} on the corresponding Cauchy slice Σ of the boundary spacetime \mathcal{B} is computed by the area of a certain bulk minimal surface which lies on $\tilde{\Sigma}$. In this case we have a lot of confidence in this entry to the AdS/CFT dictionary; firstly the RT formula obeys rather non-trivial general properties of entanglement entropies such as strong subadditivity [8–10], and secondly a general argument has been given for it in the context of Euclidean quantum gravity [11].

However, it should be clear from the outset that restricting oneself to static states is overly limiting. Not only is the field theory notion of entanglement entropy valid in a broader, time-dependent, context, but more importantly, one cannot hope to infer all possible constraints on the holographic map without considering time dependence.

The HRT proposal, which generalizes the RT construction to arbitrary time-dependent configurations by promoting a minimal surface on $\tilde{\Sigma}$ to an *extremal* surface $\mathcal{E}_{\mathcal{A}}$ in \mathcal{M} , allows one to confront geometric questions in complete generality. However, this proposal has passed far fewer checks, and an argument deriving it from first principles is still lacking. This presents a compelling opportunity to test the construction against field-theory expectations and see how it holds up. Since the new ingredient in HRT is time-dependence, the crucial property to check is causality. The present discussion therefore focuses on verifying that *the HRT prescription is consistent with field-theory causality*.²

Let us start by considering the implications of CFT causality on entanglement entropy, in order to extract the corresponding requirements to be upheld by its putative bulk dual. As we will explain in detail in §2, there are two such requirements. First, the entanglement entropy is a so-called *wedge observable*. This means that two spatial regions \mathcal{A} , \mathcal{A}' that share the same domain of dependence, $D[\mathcal{A}] = D[\mathcal{A}']$, have the same entanglement entropy, $S_{\mathcal{A}} = S_{\mathcal{A}'}$; this follows from the fact that the corresponding

¹ We will focus exclusively on local QFTs with conformal UV fixed points which are holographically dual to asymptotically AdS spacetimes in *two-derivative* theories of gravity.

² As we elaborate in the course of our discussion this result follows from Theorem 6 of [12]. As this is however not widely appreciated we focus on proving the result from a different perspective highlighting certain novel bulk constructs in the process.

reduced density matrices $\rho_{\mathcal{A}}, \rho_{\mathcal{A}'}$ are unitarily related [13]. Second, fixing the initial state, a perturbation to the Hamiltonian with support contained entirely inside $D[\mathcal{A}] \cup D[\mathcal{A}^c]$ (where \mathcal{A}^c is the complement of \mathcal{A} on a Cauchy slice) cannot affect $S_{\mathcal{A}}$. The reason is that we can choose a Cauchy slice Σ' that lies to the past of the support and contains a region \mathcal{A}' with $D[\mathcal{A}'] = D[\mathcal{A}]$; since the perturbation cannot change the state on Σ' , it cannot affect $S_{\mathcal{A}'}$, which by the previous requirement equals $S_{\mathcal{A}}$. Time-reversing the argument shows that, similarly, $S_{\mathcal{A}}$ cannot be affected by a perturbation in $D[\mathcal{A}] \cup D[\mathcal{A}^c]$ when we consider time evolution toward the past with a fixed final state.

Having specified the implications of causality for the entanglement entropy in the field theory, let us now translate them into requirements on its holographic dual. First, in order to ensure that the HRT formula in general gives the same entanglement entropy for \mathcal{A} and \mathcal{A}' , they should have the same extremal surface, $\mathcal{E}_{\mathcal{A}} = \mathcal{E}_{\mathcal{A}'}$. Second, in order for $\mathcal{E}_{\mathcal{A}}$ to be safe from influence by perturbations of the boundary Hamiltonian in $D[\mathcal{A}]$ and $D[\mathcal{A}^c]$ (when evolving either toward the future or toward the past), it has to be causally disconnected from those two regions. This means that the extremal surface has to lie in a region which we dub the *causal shadow*, denoted by $\mathcal{Q}_{\partial\mathcal{A}}$ and defined in (2.7) as the set of bulk points which are spacelike-separated from $D[\mathcal{A}] \cup D[\mathcal{A}^c]$.

This causality requirement takes an interesting guise in the case where \mathcal{A} is an entire Cauchy slice for a boundary. If this is the only boundary, and the bulk is causally trivial, then there is no causal shadow; indeed, $\mathcal{E}_{\mathcal{A}} = \emptyset$, corresponding to the fact that the entanglement entropy of the full system vanishes in a pure state. However, if the state is *not* pure, the bulk geometry is causally nontrivial: typically the bulk black-hole spacetime has two boundaries, dual to two field theories in an entangled state (which can be thought of as purifying the thermal state of the theory on one boundary). If we take the region \mathcal{A} to be a Cauchy slice for one boundary and \mathcal{A}^c a Cauchy slice for the other, then the extremal surface whose area, according to HRT, measures the amount of entanglement between the two field theories must lie in a region out of causal contact with either boundary.³

How trivial or expected is the claim that the extremal surface resides in the causal shadow? It is interesting to note that for *local* CFT observables, analogous causality violation is in fact disallowed by the gravitational time-delay theorem of Gao and Wald [14]. This theorem, which assumes that the bulk satisfies the null energy condition, implies that a signal from one boundary point to another cannot propagate faster through the bulk than along the boundary, ensuring that bulk causality respects

³ For the well-known eternal static Schwarzschild-AdS case, the shadow region degenerates to the bifurcation surface, but we will see that in general it is a finite codimension-zero bulk region.




Fig. 1: For AdS_3 , the RT formula satisfies field-theory causality marginally. The plane generated by null geodesics (color-coded by angular momentum) from a given boundary point (blue) is also ruled by spacelike geodesics at constant time (color-coded by time).

boundary causality. However, since entanglement entropy is a more nonlocal quantity, which according to HRT is captured by a bulk surface that can go behind event and apparent horizons [15, 16] and penetrate into causally disconnected regions from the boundary, it is far less obvious whether CFT causality will survive in this context.

Let us first consider a static example. Although it is guaranteed to be consistent with CFT causality since it is covered by the RT prescription which is “derived” from first principles, it is useful to gain appreciation for how innocuous or far-fetched causality violation would appear in the more general case. Intriguingly, already the simplest case of pure AdS reveals the potential for things to go wrong. As illustrated in Fig. 1, the null congruence from a single boundary point (which bounds the bulk region which a boundary source at that point can influence) is simultaneously foliated by spacelike geodesics $\{\mathcal{E}_{\mathcal{A}}\}$. So a signal that can influence a given extremal surface $\mathcal{E}_{\mathcal{A}}$ in that set can also influence $\partial\mathcal{A}$, thereby upholding CFT causality. However, note that here causality was maintained marginally: if the extremal surface was deformed away from \mathcal{A} by arbitrarily small amount, one would immediately be in danger of causality violation.

Another, less trivial, test case is the static eternal Schwarzschild-AdS black hole. The extremal surface that encodes entanglement between the two boundaries is the horizon bifurcation surface. Again, arbitrarily small deformation of this surface would shift it into causal contact with at least one of the boundaries, thereby endangering causality; in particular, entanglement entropy for one CFT should not be influenced by deformations in the other CFT. For static geometries we’re in fact safe because

extremal surfaces do not penetrate event horizons [17]; however this is no longer the case in dynamical situations [15, 16, 18–20]. Moreover, as illustrated in [21], in Vaidya-AdS geometry, $\mathcal{E}_{\mathcal{A}}$ can be null-related to the past tip of $D[\mathcal{A}]$, thereby again upholding causality just marginally—an arbitrarily small outward deformation of the extremal surface would render it causally accessible from $D[\mathcal{A}]$. These considerations demonstrate that the question of whether the HRT prescription is consistent with field-theory causality is a highly nontrivial one.

The main result of this paper is a proof that, if the bulk spacetime metric obeys the null energy condition, then the extremal surface $\mathcal{E}_{\mathcal{A}}$ does indeed obey both of the above requirements. We conclude that the HRT formula is consistent with field-theory causality. This theorem can be viewed as a generalization of the Gao-Wald theorem [14]. We regard it as a highly nontrivial piece of evidence in favor of the HRT formula. Along the way, we will also slightly sharpen the statement of the HRT formula, and in particular clarify the homology condition on $\mathcal{E}_{\mathcal{A}}$.

Partial progress towards this result was achieved in [22, 23], which showed that the extremal surface $\mathcal{E}_{\mathcal{A}}$ generically lies outside of the “causal wedge” of $D[\mathcal{A}]$, the intersection of the bulk causal future and causal past of $D[\mathcal{A}]$. (However, these works did not make the connection to field-theory causality). A stronger statement equivalent to our theorem was proved in [12] (cf., Theorem 6) and it is noted in passing that this would ensure field theory causality. We present an alternate proof which brings out some of the bulk regions more cleanly and make the connections with boundary causality more manifest.

As a byproduct of our analysis, we will identify a certain bulk spacetime region, which we call the *entanglement wedge* and denote $\mathcal{W}_{\mathcal{E}}[\mathcal{A}]$, which is bounded on one side by $D[\mathcal{A}]$ and on the other by $\mathcal{E}_{\mathcal{A}}$. Apart from providing a useful quantity in formulating and deriving our results, the entanglement wedge is, as we will argue, the bulk region most naturally associated with the boundary reduced density matrix $\rho_{\mathcal{A}}$.

The outline of this paper is as follows. We begin in §2 with an overview of the causal domains of interest on each side of the gauge-gravity duality, and motivate and state the core theorem of the paper, which shows that the HRT proposal is consistent with boundary causality. We motivate one of the major implications of our theorem by considering spherically symmetric deformations of the eternal black hole containing a region out of causal contact with both asymptotically AdS boundaries, the *causal shadow*, and showing that the HRT surface lies in this causal shadow. In §3, we begin to develop some intuition used in the proof of our main theorem, by considering classes of null geodesic congruences in AdS₃. In §4 we prove the general theorem which establishes the main result of the paper. We conclude in §5 with a discussion of the physical implications of our result and open questions.

Note added: While this paper was nearing completion [24] appeared on the arXiv, which has some overlap with the present work. It introduces the notion of quantum extremal surfaces and argues that for bulk theories that satisfy the generalized second law such surfaces satisfy the causality constraint.

2 Causal domains and entanglement entropy

In this section we will state our basic results and discuss some of their implications. The specific proof, and some additional results, will be presented in §4. In §5 we will suggest some further interpretations of our results, particularly regarding the dual of the reduced density matrix.

We will open in §2.1 by deriving the causality properties of entanglement entropy in a QFT, and setting up some notation regarding causal domains which will be useful in the sequel. In §2.2, we will review the HRT formula and discuss various causal regions in the bulk. In §2.3, we state the basic theorem and some implications for the bulk causal structure relative to specific regions arising in the HRT conjecture. §2.4 spells out a particular consequence of our results for spacetimes with multiple boundaries.

Where left unspecified, our notation follows [25].

2.1 Causality of entanglement entropy in QFT

Consider a local quantum field theory (QFT) on a d -dimensional globally hyperbolic spacetime \mathcal{B} . The state on a given Cauchy slice⁴ Σ is described by a density matrix ρ_Σ ; this could be a pure or mixed state. We are interested in the entanglement between the degrees of freedom in a region⁵ $\mathcal{A} \subset \Sigma$ and its complement \mathcal{A}^c . Following established terminology, we call the boundary $\partial\mathcal{A}$ the *entangling surface*.

The entanglement entropy is defined by first decomposing the Hilbert space \mathcal{H} of the QFT into $\mathcal{H}_\mathcal{A} \otimes \mathcal{H}_{\mathcal{A}^c}$, after imposing some suitable cutoff.⁶ The reduced density matrix $\rho_\mathcal{A} := \text{Tr}_{\mathcal{H}_{\mathcal{A}^c}} \rho_\Sigma$ captures the entanglement between \mathcal{A} and \mathcal{A}^c ; in particular, the entanglement entropy is given by its von Neumann entropy: $S_\mathcal{A} := -\text{Tr}(\rho_\mathcal{A} \ln \rho_\mathcal{A})$. For

⁴ Throughout this paper we will require all Cauchy slices to be *acausal* (no two points are connected by a causal curve). This is slightly different from the standard definition in the general-relativity literature, in which a Cauchy slice is merely required to be achronal. The reason is to ensure that different points represent independent degrees of freedom, which is useful when we decompose the Hilbert space according to subsets of the Cauchy slice.

⁵ Technically, \mathcal{A} is defined as the interior of a codimension-zero submanifold-with-boundary in Σ , $\partial\mathcal{A}$ is the boundary of that submanifold, and $\mathcal{A}^c := \Sigma \setminus (\mathcal{A} \cup \partial\mathcal{A})$.

⁶ In the case of gauge fields, this decomposition is not possible even on the lattice. Instead, one must extend the Hilbert spaces $\mathcal{H}_\mathcal{A}$, $\mathcal{H}_{\mathcal{A}^c}$ to each include degrees of freedom on $\partial\mathcal{A}$, so that $\mathcal{H} \subset \mathcal{H}_\mathcal{A} \otimes \mathcal{H}_{\mathcal{A}^c}$ [26–29].

holographic theories, we expect that this quantity has good properties in the large- N limit,⁷ unlike the Rényi entropies $S_{n,\mathcal{A}} := -\frac{1}{n-1} \ln \text{Tr}(\rho_{\mathcal{A}}^n)$ [10, 31]. Note that both quantities are determined by the eigenvalues of $\rho_{\mathcal{A}}$, and are thus insensitive to unitary transformations of $\rho_{\mathcal{A}}$.

Now, since Σ is a Cauchy slice, the future (past) evolution of initial data on it allows us to reconstruct the state of the QFT on the entirety of \mathcal{B} . In other words, the past and future domains of dependence of Σ , $D^{\pm}[\Sigma]$, together make up the background spacetime on which the QFT lives, i.e., $D^+[\Sigma] \cup D^-[\Sigma] = \mathcal{B}$. Likewise, the domain of dependence of \mathcal{A} , $D[\mathcal{A}] = D^+[\mathcal{A}] \cup D^-[\mathcal{A}]$, is the region where the reduced density matrix $\rho_{\mathcal{A}}$ can be uniquely evolved once we know the Hamiltonian acting on the reduced system in \mathcal{A} .⁸

\mathcal{A}^c similarly has its domain of dependence $D[\mathcal{A}^c]$. However, unless \mathcal{A} comprises the entire Cauchy slice, the two domains do not make up the full spacetime, $D[\mathcal{A}] \cup D[\mathcal{A}^c] \neq \mathcal{B}$, since we have to account for the regions which can be influenced by the entangling surface $\partial\mathcal{A}$. Denoting the causal future (past) of a point $p \in \mathcal{B}$ by $J^{\pm}(p)$ we find that we have to keep track of the regions $J^{\pm}[\partial\mathcal{A}]$ which are not contained in either $D[\mathcal{A}]$ or $D[\mathcal{A}^c]$. As a result, the full spacetime \mathcal{B} decomposes into four causally-defined regions: the domains of dependence of the region and its complement, and the causal future and past of the entangling surface:

$$\mathcal{B} = D[\mathcal{A}] \cup D[\mathcal{A}^c] \cup J^+[\partial\mathcal{A}] \cup J^-[\partial\mathcal{A}]. \quad (2.1)$$

These four regions are non-overlapping (except that $J^{\pm}[\partial\mathcal{A}]$ both include $\partial\mathcal{A}$). See Fig. 2 for an illustration of this decomposition. Although this decomposition is fairly obvious pictorially, for completeness we provide a proof in §4 (cf. theorem 12).

The decomposition (2.1) is particularly convenient for formulating the QFT causality constraint. Recall that the eigenvalues of the reduced density matrix $\rho_{\mathcal{A}}$, and hence the Rényi and von Neumann entropies, are invariant under unitary transformations which act on $\mathcal{H}_{\mathcal{A}}$ alone or on $\mathcal{H}_{\mathcal{A}^c}$ alone. These include perturbations of the Hamiltonian and local unitary transformations supported in the domains $D[\mathcal{A}]$ or $D[\mathcal{A}^c]$. In particular, if we consider another region \mathcal{A}' of a Cauchy slice Σ' such that $D[\mathcal{A}] = D[\mathcal{A}^c]$ (as indicated in Fig. 2), then the state $\rho_{\Sigma'}$ is related by a unitary transformation to the state ρ_{Σ} . It is clear that such a transformation can be constructed from operators localized in \mathcal{A} , and so does not change the entanglement spectrum of $\rho_{\mathcal{A}}$. Furthermore,

⁷ Technically, by “large- N ” we mean large c_{eff} , where c_{eff} is a general count of the degrees of freedom (see [30] for the general definition of c_{eff}).

⁸ We remind the reader that $D[\mathcal{A}]$ is defined as the set of points in \mathcal{B} through which every inextendible causal curve intersects \mathcal{A} . Note that, given that we have defined \mathcal{A} as an open subset of Σ , $D[\mathcal{A}]$ is open subset of \mathcal{B} .




Fig. 2: An illustration of the causal domains associated with a region \mathcal{A} , making manifest the decomposition of the spacetime into the four distinct domains indicated in (2.1). Two deformations \mathcal{A}' are also included for illustration in the right panel.

if we fix the state at $t \rightarrow -\infty$, then a perturbation to the Hamiltonian with support R cannot affect the state on a Cauchy slice to the past of R (i.e. that doesn't intersect $J^+[R]$). Such a perturbation can therefore affect the entanglement spectrum only if R intersects $J^-[\partial\mathcal{A}]$, because otherwise we can imagine evaluating $S_{\mathcal{A}}$ by using a sufficiently early Cauchy slice $\Sigma' \supset \partial\mathcal{A}$ that passes to the past of R . Similarly, if we fix the state at $t \rightarrow +\infty$, the spectrum can be affected only by perturbations in $J^+[\partial\mathcal{A}]$. In summary, we have the following properties of $\rho_{\mathcal{A}}$:

- The entanglement spectrum of $\rho_{\mathcal{A}}$ depends only on the domain $D[\mathcal{A}]$ and not on the particular choice of Cauchy slice Σ . The spectrum is thus a so-called “wedge observable” (although it is not, of course, an observable in the usual sense).
- Fixing the state in either the far past or the far future, the entanglement spectrum of $\rho_{\mathcal{A}}$ is insensitive to any local deformations of the Hamiltonian in $D[\mathcal{A}]$ or $D[\mathcal{A}^c]$.

These are the crucial causality requirements that entanglement (Rényi) entropies are required to satisfy in any relativistic QFT.

The essential result of this paper is that the HRT proposal for computing $S_{\mathcal{A}}$ satisfies these causality constraints. In the conclusions we will revisit the question of what the dual of $\rho_{\mathcal{A}}$, and thus of the data in $D[\mathcal{A}]$, might be.

2.2 Bulk geometry and holographic entanglement entropy

Let us now restrict attention to the class of holographic QFTs, which are theories dual to classical dynamics in some bulk asymptotically AdS spacetime. To be precise, we only consider strongly coupled QFTs in which the classical gravitational dynamics

truncates to that of Einstein gravity, possibly coupled to matter which we will assume satisfies the null energy condition.

The dynamics of the QFT on \mathcal{B} is described by classical gravitational dynamics on a bulk asymptotically locally AdS spacetime \mathcal{M} with conformal boundary \mathcal{B} , the spacetime where the field theory lives. We define $\tilde{\mathcal{M}} := \mathcal{M} \cup \mathcal{B}$. $\tilde{\mathcal{M}}$ is endowed with a metric \tilde{g}_{ab} which is related by a Weyl transformation to the physical metric g_{ab} on \mathcal{M} , $\tilde{g}_{ab} = \Omega^2 g_{ab}$, where $\Omega \rightarrow 0$ on \mathcal{B} .⁹ Causal domains on $\tilde{\mathcal{M}}$ will be denoted with a tilde to distinguish them from their boundary counterparts, e.g., $\tilde{J}^\pm(p)$ will denote the causal future and past of a point p in $\tilde{\mathcal{M}}$ and $\tilde{D}[R]$ will denote the domain of dependence of some set $R \subset \tilde{\mathcal{M}}$.

It will also be useful to introduce a compact notation to indicate when two points p and q are spacelike-separated; for this we adopt the notation \asymp , i.e.

$$p \asymp q \Leftrightarrow \nexists \text{ a causal curve between } p \text{ and } q. \quad (2.2)$$

Moreover, to denote regions that are spacelike separated from a point, we will use $\mathcal{S}(p)$ and $\tilde{\mathcal{S}}(p)$ in the boundary and bulk respectively,

$$\mathcal{S}(p) := \{q \mid p \asymp q\} = (J^+(p) \cup J^-(p))^c \quad \text{and} \quad \tilde{\mathcal{S}}(p) := \left(\tilde{J}^+(p) \cup \tilde{J}^-(p)\right)^c. \quad (2.3)$$

Just as for other causal sets, we can extend these definitions to any region R , namely $\mathcal{S}[R] := \cap_{p \in R} \mathcal{S}(p)$ is the set of points which are causally disconnected from the entire region R , etc.

Having established our notation for general causal relations, let us now specify the notation relevant for holographic entanglement entropy. As before we will fix a region \mathcal{A} on the boundary. The HRT proposal [3] states that the entanglement entropy $S_{\mathcal{A}}$ is holographically computed by the area of a bulk codimension-two extremal surface $\mathcal{E}_{\mathcal{A}}$ that is anchored on $\partial\mathcal{A}$; specifically,

$$S_{\mathcal{A}} = \frac{\text{Area}(\mathcal{E}_{\mathcal{A}})}{4G_N}. \quad (2.4)$$

In the static (RT) case, it is known that the extremal surface is required to be homologous to \mathcal{A} , meaning that there exists a bulk region $\mathcal{R}_{\mathcal{A}}$ such that $\partial\mathcal{R}_{\mathcal{A}} = \mathcal{A} \cup \mathcal{E}_{\mathcal{A}}$. So far, it has not been entirely clear what the correct covariant generalization of this condition is. In particular, should it merely be a topological condition, or should one impose geometrical or causal requirements on $\mathcal{R}_{\mathcal{A}}$, for example, that it be spacelike? (A critical discussion of the issues involved can be found in [32].) In this paper, we

⁹ These are necessary but not sufficient conditions for the spacetime to be asymptotically AdS.

will show that a clean picture, consistent with all aspects of field-theory causality, is obtained by requiring that $\mathcal{R}_{\mathcal{A}}$ be a region of a bulk Cauchy slice.¹⁰ We will call this the “spacelike homology” condition.¹¹

The homology surface $\mathcal{R}_{\mathcal{A}}$ naturally leads us to the key construct pertaining to entanglement entropy, which we call the *entanglement wedge* of \mathcal{A} , denoted by¹² $\mathcal{W}_{\mathcal{E}}[\mathcal{A}]$. This can be defined as a causal set, namely the bulk domain of dependence of $\mathcal{R}_{\mathcal{A}}$,

$$\mathcal{W}_{\mathcal{E}}[\mathcal{A}] := \tilde{D}[\mathcal{R}_{\mathcal{A}}]. \quad (2.5)$$

Note that the entanglement wedge is a bulk codimension-zero spacetime region, which can be equivalently identified with the region defined by the set of bulk points which are spacelike-separated from $\mathcal{E}_{\mathcal{A}}$ and connected to $D[\mathcal{A}]$. The latter definition has the advantage of absolving us of having to specify an arbitrary homology surface $\mathcal{R}_{\mathcal{A}}$ rather than just $\mathcal{E}_{\mathcal{A}}$ and $D[\mathcal{A}]$. As we shall see below, the bulk spacetime can be naturally decomposed into four regions analogously to the boundary decomposition (2.1); the entanglement wedge is then the region associated with (and ending on) $D[\mathcal{A}]$.

While we have focused on the regions in the bulk which enter the holographic entanglement entropy constructions, we pause here to note two other causal constructs that can be naturally associated with \mathcal{A} . First of all we have the *causal wedge* $\mathcal{W}_{\mathcal{C}}[\mathcal{A}]$ which is set of all bulk points which can both send signals to and receive signals from boundary points contained in $D[\mathcal{A}]$, i.e.,¹³

$$\mathcal{W}_{\mathcal{C}}[\mathcal{A}] := \tilde{J}^+[D[\mathcal{A}]] \cap \tilde{J}^-[D[\mathcal{A}]]. \quad (2.6)$$

(The entanglement wedge $\mathcal{W}_{\mathcal{E}}[\mathcal{A}]$ and causal wedge $\mathcal{W}_{\mathcal{C}}[\mathcal{A}]$ are in fact special cases of the “rim wedge” and “strip wedge” introduced recently in [33] as bulk regions associated with residual entropy.)

The second bulk causal domain which will play a major role in our discussion below is a region we call the *causal shadow* $\mathcal{Q}_{\partial\mathcal{A}}$ associated with the entangling surface $\partial\mathcal{A}$.

¹⁰ Technically, similarly to \mathcal{A} , we define $\mathcal{R}_{\mathcal{A}}$ to be the interior of a codimension-zero submanifold-with-boundary of a Cauchy slice $\tilde{\Sigma}$ of $\tilde{\mathcal{M}}$ (with $\tilde{\Sigma} \cap \mathcal{B} = \Sigma$). Since $\tilde{\Sigma}$ itself has a boundary (namely its intersection with \mathcal{B}), the interior of a subset (in the sense of point-set topology) includes the part of its boundary along \mathcal{B} . Thus, $\mathcal{R}_{\mathcal{A}}$ includes \mathcal{A} (but not $\mathcal{E}_{\mathcal{A}}$).

¹¹ If there are multiple extremal surfaces obeying the spacelike homology condition, then we are to pick the one with smallest area. However, in this paper we will not use this additional minimality requirement; all our theorems apply to any spacelike-homologous extremal surface.

¹² While we have associated it notationally with the region \mathcal{A} , it depends only on $D[\mathcal{A}]$.

¹³ Following [22], we can also define a particular bulk codimension-two surface $\Xi_{\mathcal{A}}$, the causal information surface, to be the rim of the causal wedge; in fact, it is the minimal area codimension-two surface lying on $\partial\mathcal{W}_{\mathcal{C}}[\mathcal{A}]$.




Fig. 3: Example of a causally trivial spacetime and a boundary region \mathcal{A} whose causal shadow is a finite spacetime region. We have engineered an asymptotically AdS_3 geometry sourced by matter satisfying the null energy condition (see footnote 14) and taken \mathcal{A} to nearly half the boundary, $\varphi_{\mathcal{A}} = 1.503$, at $t = 0$ (thick red curve). The shaded regions on the boundary cylinder are $D[\mathcal{A}]$ and $D[\mathcal{A}^c]$ respectively. The extremal surface is the thick blue curve, while the purple curves are the rims of the causal wedge (causal information surfaces) for \mathcal{A} and \mathcal{A}^c respectively. A few representative generators are provided for orientation: the blue null geodesics generate the boundary of the causal wedge for \mathcal{A} while the green ones do likewise for \mathcal{A}^c . The orange generators in the middle of the spacetime generate the boundary of the causal shadow region $\mathcal{Q}_{\partial\mathcal{A}}$.

We define this region as the set of points in the bulk \mathcal{M} that are spacelike-related to both $D[\mathcal{A}]$ and $D[\mathcal{A}^c]$, i.e.,

$$\begin{aligned} \mathcal{Q}_{\partial\mathcal{A}} &:= \left(\tilde{J}^+[D[\mathcal{A}]] \cup \tilde{J}^-[D[\mathcal{A}]] \cup \tilde{J}^+[D[\mathcal{A}^c]] \cup \tilde{J}^-[D[\mathcal{A}^c]] \right)^c \\ &= \tilde{\mathcal{S}}[D[\mathcal{A}] \cup D[\mathcal{A}^c]]. \end{aligned} \quad (2.7)$$

For a generic region \mathcal{A} in a generic asymptotically AdS spacetime, the causal shadow

is a codimension-zero spacetime region; see Fig. 3 for an illustrative example.¹⁴ In certain special (but familiar) situations, such as spherically symmetric regions in pure AdS (where $\rho_{\mathcal{A}}$ is unitarily equivalent to a thermal density matrix), it can degenerate to a codimension-two surface. In such special cases, the entanglement wedge and the causal wedge coincide [22]. In general, the causal information surface for \mathcal{A} and that for \mathcal{A}^c comprise the edges of the causal shadow. For a generic pure state these causal information surfaces each recede from $\mathcal{E}_{\mathcal{A}}$ towards their respective boundary region but approach each other near the AdS boundary. Hence the geometrical structure of $\mathcal{Q}_{\partial\mathcal{A}}$, described in language of a three-dimensional bulk, is a “tube” (connecting the two components of $\partial\mathcal{A}$) with a diamond cross-section, which shrinks to a point where the tube meets the AdS boundary at $\partial\mathcal{A}$.

For topologically trivial deformations of AdS, in the absence of $\mathcal{E}_{\mathcal{A}}$ (i.e. when the state is pure and $\mathcal{A} = \Sigma$) the causal shadow disappears, but intriguingly, even when \mathcal{A} is the entire boundary Cauchy slice, the causal shadow can be nontrivial. This occurs for example in the AdS₃-geon spacetimes¹⁵ [34] and in perturbations of the eternal AdS black hole, such as those studied by [35]. In such a situation we simply define the causal shadow of the entire boundary (dropping the subscript) as

$$\mathcal{Q} := \tilde{\mathcal{S}}[\mathcal{B}] = \left(\tilde{J}^+[\mathcal{B}] \cup \tilde{J}^-[\mathcal{B}] \right)^c \quad (2.8)$$

Here \mathcal{B} is understood generally to include multiple disconnected components; the causal shadow is the region spacelike separated from points on all the boundaries.

2.3 Causality constraints on extremal surfaces

Having developed the various causal concepts which we require, let us now ask what the constraints of field-theory causality concerning entanglement entropy translate to in the bulk. The first constraint is that $S_{\mathcal{A}}$ should be a wedge observable, i.e. if $D[\mathcal{A}] = D[\mathcal{A}']$ then $S_{\mathcal{A}} = S_{\mathcal{A}'}$. For this to hold in general, we need $\mathcal{E}_{\mathcal{A}} = \mathcal{E}_{\mathcal{A}'}$. The second concerns perturbations of the field-theory Hamiltonian. Such perturbations will source perturbations of the bulk fields, including the metric, that will travel causally with respect to the background metric. In particular, disturbances originating in $D[\mathcal{A}]$

¹⁴ The bulk metric used in the plot for Fig. 3 is

$$ds^2 = \frac{1}{\cos^2 \rho} \left(-f(\rho) dt^2 + \frac{d\rho^2}{f(\rho)} + \sin^2 \rho d\varphi^2 \right), \quad f(\rho) = 1 - \frac{1}{2} \sin^2(2\rho).$$

The matter supporting this geometry satisfies the null energy condition as can be checked explicitly.

¹⁵ Since these describe pure states, the presence of a causal shadow region does not necessarily guarantee the presence of an extremal surface *whose area gives the entanglement entropy* contained within it. However, there will be *some* extremal surface spanning this region.

will be dual to bulk modes propagating in $\tilde{J}^+[D[\mathcal{A}]]$ (if we fix the state in the far past) or in $\tilde{J}^-[D[\mathcal{A}]]$ (if we fix the state in the far future). If either of these bulk regions intersected $\mathcal{E}_{\mathcal{A}}$, the dual of local operator insertions in $D[\mathcal{A}]$ could change the area of $\mathcal{E}_{\mathcal{A}}$, meaning that the HRT proposal would be inconsistent with causality in the QFT. By the same token, the extremal surface cannot intersect $\tilde{J}^+[D[\mathcal{A}^c]]$ or $\tilde{J}^-[D[\mathcal{A}^c]]$. Since the region complement to union of the causal sets $\tilde{J}^\pm[D[\mathcal{A}]]$, $\tilde{J}^\pm[D[\mathcal{A}^c]]$ is the set of points that are spacelike related to $D[\mathcal{A}] \cup D[\mathcal{A}^c]$, we learn that

$$\mathcal{E}_{\mathcal{A}} \simeq D[\mathcal{A}] \cup D[\mathcal{A}^c]. \quad (2.9)$$

In others words, using (2.7) we can say that $\mathcal{E}_{\mathcal{A}}$ has to lie in the causal shadow of $\partial\mathcal{A}$

$$\mathcal{E}_{\mathcal{A}} \subset \mathcal{Q}_{\partial\mathcal{A}}. \quad (2.10)$$

It is known, based on properties of extremal surfaces, that $\mathcal{E}_{\mathcal{A}}$ lies outside the causal wedges $\mathcal{W}_c[\mathcal{A}]$ and $\mathcal{W}_c[\mathcal{A}^c]$ [12, 22, 23]. This leaves open the possibility that the surface could still lie in the causal future (or past) of the boundary domain of dependence of \mathcal{A} or \mathcal{A}^c . A particular worry arises in explicit examples in Vaidya-AdS geometries where the extremal surface lies on the boundary of $\tilde{J}^+[D[\mathcal{A}]]$. This then leaves open the question whether one might indeed be able to push $\mathcal{E}_{\mathcal{A}}$ into a causally forbidden region, by introducing appropriate deformations in $D[\mathcal{A}]$. A theorem of Wall [12] (Theorem 6 of the reference), guarantees that this does not occur (modulo some assumptions).

We will prove an essentially equivalent statement in §4, directly for extremal surfaces in an asymptotically AdS spacetime. The main result however can be stated in terms of three simple causal relations:

$$\begin{aligned} \tilde{D}[\mathcal{R}_{\mathcal{A}}] \cap \mathcal{B} &= D[\mathcal{A}] \\ \tilde{D}[\mathcal{R}_{\mathcal{A}^c}] \cap \mathcal{B} &= D[\mathcal{A}^c] \\ \tilde{J}^\pm[\mathcal{E}_{\mathcal{A}}] \cap \mathcal{B} &= J^\pm[\partial\mathcal{A}]. \end{aligned} \quad (2.11)$$

In other words, the causal split of the bulk into spacelike- and timelike-separated regions from $\mathcal{E}_{\mathcal{A}}$ restricts to the boundary at precisely the boundary split (2.1). Given the decomposition (2.1), these causal relations imply that perturbations in $D[\mathcal{A}] \cup D[\mathcal{A}^c]$ are not in causal contact with $\mathcal{E}_{\mathcal{A}}$. So, as required, the extremal surface lies in the causal shadow.

As a consequence of this theorem, we will also show that, if there is a spacelike region \mathcal{A}' such that $D[\mathcal{A}'] = D[\mathcal{A}]$, then there is a bulk region $\mathcal{R}_{\mathcal{A}'}$ such that $\partial\mathcal{R}_{\mathcal{A}'} = \mathcal{A}' \cup \mathcal{E}_{\mathcal{A}}$, so $\mathcal{E}_{\mathcal{A}}$ is spacelike-homologous to \mathcal{A}' . Thus, the HRT formula gives the same entanglement entropy for \mathcal{A}' and \mathcal{A} , as required on the field-theory side.

2.4 Entanglement for disconnected boundary regions

A striking consequence of the theorems discussed above emerges when we consider spacetimes with two boundary components, and let \mathcal{A} be (a Cauchy slice for) all of one component.

As a starting point, consider the eternal Schwarzschild-AdS $_{d+1}$ black hole in the Hartle-Hawking state, with a Penrose diagram shown in Fig. 4(a) below. The left and right boundaries of the diagram each have the topology $\mathbf{S}^{d-1} \times \mathbb{R}$. This geometry is believed to be dual to the CFT on the product spatial geometry $\mathbf{S}_L^{d-1} \times \mathbf{S}_R^{d-1}$, in the entangled “thermofield double” state [36–39]:

$$|\text{HH}\rangle_{L,R} = \sum_i e^{-\frac{1}{2}\beta E_i} |E_i\rangle_L |E_i\rangle_R \quad (2.12)$$

where $|E_i\rangle_{R,L}$ is the energy eigenstate of the CFT on $\mathbf{S}_{R,L}^{d-1}$.

Let Σ_R lie on the $t = 0$ slice of the right boundary, and consider the reduced density matrix for some region $\mathcal{A} \subset \Sigma_R$. Since this is a static geometry, its entanglement entropy $S_{\mathcal{A}}$ is computed by a minimal surface $\mathcal{E}_{\mathcal{A}}$ which never penetrates past the bifurcation surface \mathcal{X} of the black hole [17].¹⁶ If we let \mathcal{A} be the full Cauchy slice of one of the boundaries, say $\mathcal{A} = \Sigma_R$, the extremal surface precisely coincides with the black hole bifurcation surface, as indicated in Fig. 4. Note that $\mathcal{E}_{\mathcal{A}}$ lies on the edge of the causally acceptable region since \mathcal{X} sits at the boundary of both $\mathcal{W}_{\mathcal{C}}[\mathcal{A}]$ and $\mathcal{W}_{\mathcal{C}}[\mathcal{A}^c]$, and therefore constitutes the entire causal shadow for this special case.

One might now wonder what happens if we deform the state (2.12). This is not an innocuous question. In time-dependent geometries, the global (teleological) nature of the event horizon implies that extremal surfaces anchored on the boundary *can* pass through this horizon [15]. Furthermore, as first explicitly shown in [16], even apparent horizons do not form a barrier to the extremal surfaces. Hence we see that, a priori, in a state which is a deformation of (2.12), $\mathcal{E}_{\mathcal{A}}$ is in danger of entering $\mathcal{W}_{\mathcal{C}}[\mathcal{A}^c]$.

The theorems we have stated above indicate that this does not happen. The question is, how precisely does the extremal surface $\mathcal{E}_{\mathcal{A}}$ avoid doing so? As a first step to answering this, consider a deformation of the static eternal case localized along a null shell emitted from the right boundary at some time. The corresponding metric is given by the global Vaidya-SAdS geometry, where both the initial (prior to the shell) and final (after the shell) spacetime regions describe a black hole. Fig. 4b presents a sketch of the Penrose diagram of such a geometry, contrasted with the standard static eternal Schwarzschild-AdS black hole (Fig. 4a). The diagonal brown line represents the shell

¹⁶ Note that the extremal surface does not come arbitrarily close to the horizon—it either includes a component that wraps the horizon, or stays a finite distance away from it [32].




Fig. 4: Sketch of Penrose diagram for (a) static eternal Schwarzschild-AdS and (b) ‘thin shell’ Vaidya-Schwarzschild-AdS, with the various regions labeled. The AdS boundaries are represented by vertical black lines, the singularities by purple curves, the horizons by diagonal blue lines, and the ‘shell’ in the Vaidya case by diagonal brown line.

which is sourced at some time on the right boundary and implodes into the black hole (terminating at the future singularity), and the blue lines represent the various (future and past, left and right) event horizons. The solid parts of these lines indicate where these event horizons coincide with apparent horizons (as well as isolated horizons); the dashed parts are parts of the event horizon which are not apparent horizons.

In such a geometry, let us again consider $\mathcal{A} = \Sigma_R$. Then our theorems guarantee that the extremal surface must lie on the null sheet separating regions R_c and P_c : it is again spacelike-separated from both $D[\Sigma_L]$ and $D[\Sigma_R]$. (In fact, since the spacetime prior to the shell is identical to the eternal static case, the extremal surface remains in the same location as for the static case, namely the bifurcation surface where regions R_c and L touch.) The situation is again marginal, much like the original undeformed case. Indeed, any perturbation to Schwarzschild-AdS which emanates from (or reaches to) the right boundary cannot change the location of the original extremal surface by causality; it could at most generate a new extremal surface.

A less marginal case occurs when we symmetrically perturb both copies of the CFT as above. Consider a perturbation at $t = 0$ such that spherically symmetric null shells are emitted both to the past and future on both sides of the diagram. One then obtains the Penrose diagram shown in Fig. 5; this has time-reflection symmetry about

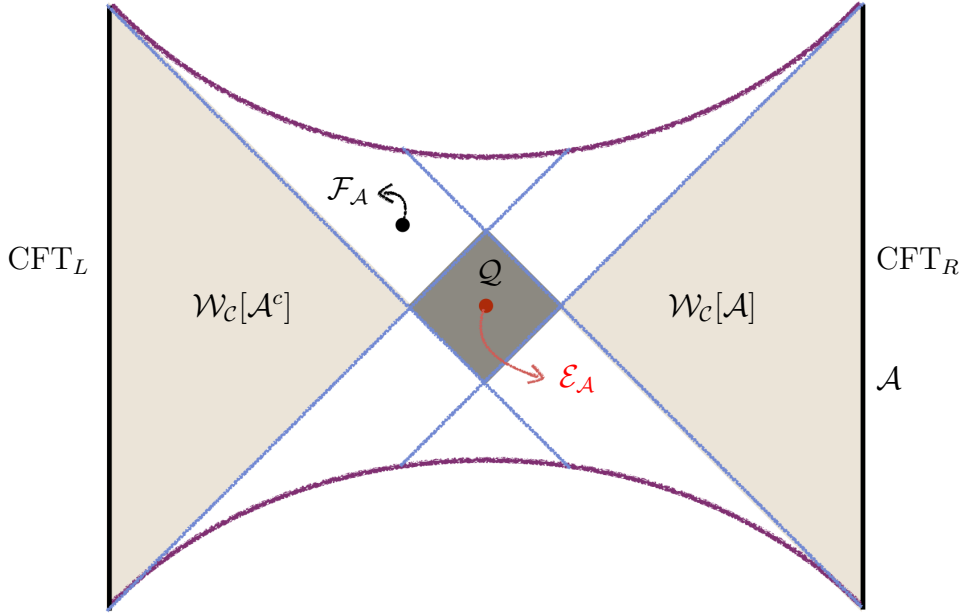


Fig. 5: Sketch of Penrose diagram for a symmetric Vaidya-Schwarzschild-AdS geometry obtained by imploding null shells to the past and future from both boundaries. The crucial new feature of note is the presence a *causal shadow* region that is spacelike separated from both boundaries. We have also indicated the extremal surface \mathcal{E}_A for the region $\mathcal{A} = \Sigma_R$ in red at the center of the figure and \mathcal{F}_A is a \mathbf{S}^{d-1} of finite area in the causal future of the left boundary. The lightly shaded regions are the causal wedges associated with \mathcal{A} and \mathcal{A}^c respectively.

$t = 0$, symmetry under exchanging the left and right sides, and the $SO(d)$ rotational symmetry.

According to the theorems above, the extremal surface must be spacelike-separated from both boundaries, when we take $\mathcal{A} = \Sigma_R$. Using both time and space reflection symmetry, it is clear that \mathcal{E}_A must sit in the center of the *causal shadow* \mathcal{Q} of the two boundaries, spacelike separated from both.

In the general case of spherically symmetric spacetime (even in the absence of time or space reflection symmetry) there is an easy proof of our claim that \mathcal{E}_A must lie in the causal shadow. We proceed by contradiction: suppose that a spherical extremal surface \mathcal{E}_A lies in $\tilde{J}^+[\Sigma_L]$. This means that on a Penrose diagram, it lies somewhere in the top-left region; say it is the surface \mathcal{F}_A indicated in Fig. 5 (which by rotational symmetry is a copy of \mathbf{S}^{d-1}). Let us then consider the past congruence of null normal geodesics from \mathcal{F}_A towards \mathcal{B}_L . Since we assume that \mathcal{F}_A candidate surface lies in $\tilde{J}^+[\Sigma_L]$, past-going null congruences from the surface intersect \mathcal{B}_L on a spacelike codimension-one surface.

In other words, the area of the spheres grows without bound along this past-directed congruence.

However, by definition, for an extremal surface the initial expansion is vanishing. Moreover, if the matter in the spacetime satisfies the null energy condition, then it also follows that the area along the congruence is guaranteed not to grow. Nor can the area go to zero along the congruence, since the area of the S^{d-1} represented by each point on the Penrose diagram is finite. It therefore follows that our assumption about $\mathcal{E}_{\mathcal{A}}$ penetrating $\tilde{\mathcal{J}}^+[\Sigma_L]$ must be erroneous; $\mathcal{F}_{\mathcal{A}}$ cannot be an extremal surface. Running a similar argument for the other unshaded regions in Fig. 5, we learn that the extremal surface must indeed lie in the causal shadow region, as denoted by the red surface $\mathcal{E}_{\mathcal{A}}$.

Indeed, in this particular case, the extremal surface lies at the point on the Penrose diagram where the future and past apparent horizons meet—the “apparent bifurcation surface”. The fact that it lies in the causal shadow is a consequence of the familiar fact that the apparent horizon can never be outside the event horizon, applied to both future and past horizons.

While the above result relied on the special properties of spherical symmetry (both of the spacetime and the null congruences therein), the theorems we prove in §4 will establish this in full generality.

In the next two sections we set out to prove the theorems stated in §2.3. The proof in our spherically symmetric case indicates that understanding null congruences leaving the extremal surface might play a key role. We will therefore spend some time in §3 examining null congruences emanating from bulk codimension-two surfaces in AdS_3 , in order to develop a picture of the relevant causal domains, before embarking on a general proof in §4.

3 Null geodesic congruences in AdS_3

In this section, we consider null geodesic congruences emanating from curves in AdS_3 that are anchored at the boundary. Our aim is to build some intuition about such congruences in a simple setting, since their properties will play a crucial role in the proofs in what follows. Readers familiar with the general statements are invited to skip ahead to the abstract discussion.

We work in the Poincaré patch of AdS_3 with the standard metric:

$$ds^2 = \frac{1}{z^2} (-dt^2 + dx^2 + dz^2) \quad (3.1)$$

Since our aim is to understand specifically the (causal) boundary of bulk causal domains, we are going to examine properties of null geodesic congruences. In particular,

for a spacelike codimension-one region $R \subset \mathcal{M}$ which is anchored on the AdS boundary, the domain of dependence $\tilde{D}[R]$ is bounded by a family of outgoing null geodesics emanating from ∂R , up to the point where each geodesic encounters a caustic or intersects another generator.¹⁷

To gain intuition for how these null congruences behave in the context of the extremal surfaces of interest, we examine a more general family of codimension-two surfaces (these are curves in AdS_3) which in the above coordinates are given by

$$x^2 + \frac{z^2}{a^2} = 1, \quad t = 0 \tag{3.2}$$

parameterized by a . Note that all of these are anchored on the boundary $\mathbb{R}^{1,1}$ at the ends of the interval $\mathcal{A} = \{(t, x) \in \mathbb{R}^{1,1} \mid t = 0, x \in [-1, 1]\}$. (For orientation, see the bottom set of curves in Fig. 7.) When $a = 1$, the surface is a semi-circle, which is simultaneously the causal information surface $\Xi_{\mathcal{A}}$ defined in [22], and the extremal surface $\mathcal{E}_{\mathcal{A}}$ for the region \mathcal{A} under consideration. Surfaces with $a < 1$ lie inside the causal wedge $\mathcal{W}_{\mathcal{C}}[\mathcal{A}]$, while those with $a > 1$ lie outside i.e., they are spacelike related to $D[\mathcal{A}]$. We wish to study the family of null congruences leaving these surfaces, as we vary a . The geodesics will be labelled by their starting position x_0 and parameterized by an affine parameter λ (fixed such that we have unit energy along each geodesic).

3.1 Explicit solutions for geodesic congruences

Since the $a = 1$ surface is extremal, the null expansion $\Theta(\lambda; a = 1) = 0$ for each generator. For the surfaces with $a < 1$, closer to the boundary, we expect that the expansion is positive and the congruence intersects the boundary in a spacelike curve inside $D[\mathcal{A}] = \{(t, x) \in \mathbb{R}^{1,1} \mid |t \pm x| \leq 1\}$. For curves with $a > 1$, long ellipse, we expect the expansion to be negative. The resulting congruence should develop a caustic before reaching the boundary.

Due to the relative simplicity of the set-up, we can confirm these expectations explicitly. Since everything is time-symmetric, let us consider just the future-directed

¹⁷ The latter set of intersections is referred to as cross-over points; the set of these generically form a crossover seam which is codimension-one on this null surface.

outgoing congruence:

$$\begin{aligned}
z(\lambda) &= \frac{a \sqrt{1-x_0^2} \sqrt{1-x_0^2+a^2x_0^2}}{a(1-x_0^2)\lambda + \sqrt{1-x_0^2+a^2x_0^2}} \\
x(\lambda) &= x_0 \frac{a(1-a^2)(1-x_0^2)\lambda + \sqrt{1-x_0^2+a^2x_0^2}}{a(1-x_0^2)\lambda + \sqrt{1-x_0^2+a^2x_0^2}} \\
t(\lambda) &= \frac{a^2(1-x_0^2)\sqrt{1-x_0^2+a^2x_0^2}\lambda}{a(1-x_0^2)\lambda + \sqrt{1-x_0^2+a^2x_0^2}}
\end{aligned} \tag{3.3}$$

Note that the endpoints of these generators at $\lambda = \infty$ are given by

$$z_\infty = 0, \quad x_\infty = x_0(1-a^2), \quad t_\infty = a\sqrt{1-x_0^2+a^2x_0^2} \tag{3.4}$$

A representative plot of the generators is given in Fig. 6 for $a = 0.5$ (left) and $a = 1.5$ (right). We see that when $a < 1$, the generators don't intersect each other before reaching the boundary, and they reach within $D^+[\mathcal{A}]$. On the other hand, when $a > 1$, the generators intersect in a seam (drawn as thick blue curve, whose explicit expression is given below in (3.5)), before reaching the boundary (with the geodesic endpoints indicated by the red curves in Fig. 6). We call the points on this seam the *cross-over points*; non-neighbouring geodesics intersect at these points. This seam terminates in a *caustic*, which as always refers to the locus where neighbouring geodesics intersect.

3.2 Intersections within congruences

We can determine the intersection between distinct geodesics in the bulk using the explicit expressions from (3.3). By symmetry of the set-up, we know that geodesics with opposite values of x_0 necessarily intersect, and they must do so at $x = x_\times = 0$. Solving for the intersection of the pair of geodesics starting from x_0 and $-x_0$ we find that they meet at:

$$t_\times = \frac{\sqrt{1-x_0^2+a^2x_0^2}}{a}, \quad z_\times = \frac{a^2-1}{a}\sqrt{1-x_0^2}, \quad \lambda_\times = \frac{\sqrt{1-x_0^2+a^2x_0^2}}{a(a^2-1)(1-x_0^2)} \tag{3.5}$$

This generates the seam of cross-over points depicted in the right panel of Fig. 6, and plotted for various values of a in Fig. 7 (the top set of curves, color-coded by a corresponding to the initial surface indicated by the thick horizontal curve of the same color). It is easy to see from (3.5) that the cross-over points terminate on the boundary at the future tip of $D^+[\mathcal{A}]$, i.e., at $z = 0$, $x = 0$, $t = 1$, corresponding to the intersection of the boundary geodesics $x_0 = \pm 1$. On the other hand, the cross-over




Fig. 6: Null normal congruence from the initial surface given by (3.2) with $a = 0.5$ (left) and $a = 1.5$ (right). The initial surface is the bold black curve on the bottom, the boundary is the shaded plane on the left in each plot (with the domain of dependence $D^+[\mathcal{A}]$ boundary indicated by the thin black lines), the individual geodesics are the thin lines color-coded by x_0 , their endpoints on the boundary are depicted by the red curve, and finally the seam of crossover points where generators intersect for $a > 1$ is the blue thick curve. (The generators are cut off at a finite value of $\lambda \approx 64$, so in the plot they don't look like they reach all the way to the boundary.)

seams for different a start at the point in the bulk when neighbouring geodesics from $x_0 \simeq 0$ intersect which happens at

$$x_{\times} = 0, \quad t_{\times} = \frac{1}{a}, \quad z_{\times} = \frac{a^2 - 1}{a}, \quad \lambda_{\times} = \frac{1}{a(a^2 - 1)} \quad (3.6)$$

To summarize, depending on whether a is greater or less than 1, the congruence has qualitatively different behaviour, as illustrated in Fig. 7. For $a < 1$ (depicted by colors from red toward green), the congruence reaches the boundary inside $D^+[\mathcal{A}]$, while for $a > 1$, the generators intersect each other at the seam of crossover points (depicted by colors from red toward purple). At precisely $a = 1$, all generators reach the boundary at the future tip of $D^+[\mathcal{A}]$, namely $z = 0$, $x = 0$, $t = 1$.




Fig. 7: Initial surfaces (thick curves at the bottom, color-coded by a), along with endpoints of the generators of the corresponding null congruence: for $a = 1$ (initial surface is the red semi-circle), all generators meet at the tip. Increasing $a > 1$ (color shift towards purple and blue) makes the generators intersect at the seam of cross-over points before reaching the boundary. On the other hand, decreasing $a < 1$ (color shift towards orange and green) makes the generators reach the boundary within $D^+[A]$ (depicted as in Fig. 6).

3.3 Expansion of congruences and caustics

Let us now analyze the expansion along this congruence. This can be calculated as the change in area along the wavefront

$$\Theta(\lambda, x_0) = \frac{1}{A(\lambda, x_0)} \frac{\partial}{\partial \lambda} A(\lambda, x_0) \quad (3.7)$$

with

$$A(\lambda, x_0) = \int_{x_0}^{x_0 + \delta x} \sqrt{\frac{-t'(\lambda, \tilde{x}_0)^2 + x'(\lambda, \tilde{x}_0)^2 + z'(\lambda, \tilde{x}_0)^2}{z^2(\lambda, \tilde{x}_0)}} d\tilde{x}_0 \quad (3.8)$$

where $t'(\lambda, x_0) \equiv \frac{\partial}{\partial x_0} t(\lambda; x_0)$ etc., using the expressions given in (3.3). While one can numerically solve for $\Theta(\lambda)$ it is easier to obtain the solution for small λ and evolve using the Raychaudhuri equation.

Near $\lambda = 0$, the leading order expression for Θ is:

$$\Theta_0 \equiv \Theta(\lambda = 0) = \frac{a(1 - a^2)(1 - x_0^2)^2}{(1 - x_0^2 + a^2 x_0^2)^{3/2}} \quad (3.9)$$

This is plotted in the left panel of Fig. 8 (with same color-coding by a as employed in Fig. 7). At the ends of the interval $x_0 = \pm 1$, Θ_0 vanishes (which is to be expected since




Fig. 8: Expansion $\Theta(\lambda; x_0)$ along the generators for various values of a (color-coded by a as in Fig. 7). On left, we show the expansion from the initial surface $\lambda = 0$ as a function of the starting position x_0 . On right, we fix $x_0 = 0$ as plot the evolution of $\Theta(\lambda)$ along the radial generator.

the congruence approximates a larger one with $a = 1$), while Θ_0 reaches its extremum at the midpoint, $x_0 = 0$ (again, expected by symmetry), where $\Theta_0(x_0 = 0) = a(1 - a^2)$. Furthermore, Θ_0 is positive for $a < 1$ and negative for $a > 1$; that is, the congruences are expanding for $a < 1$ and converging for $a > 1$). The former make it out to the boundary without intersecting, while the latter have a seam of cross-overs. As we will see below, the geodesics end in a curve of caustics, which touches the seam of cross-overs at the endpoint of the latter.

Given Θ_0 as our initial condition, it is straightforward to solve the Raychaudhuri equation

$$\frac{d\Theta}{d\lambda} = -\Theta^2 - 2\sigma_{ab}\sigma^{ab} - R_{ab}\xi^a\xi^b \quad (3.10)$$

to find the expansion along the geodesics. Here ξ^a is the tangent vector to the null geodesics and $\sigma_{\mu\nu}$ is the shear of the congruence. For a one-dimensional congruence the shear trivially vanishes and the Ricci tensor contracted with null tangents likewise vanishes upon using the bulk equations of motion $R_{ab} = -2g_{ab}$, so (3.10) simplifies to:

$$\frac{d\Theta}{d\lambda} = -\Theta^2 \quad \Rightarrow \quad \Theta(\lambda) = \frac{\Theta_0}{1 + \Theta_0\lambda} \quad (3.11)$$

Using (3.9), we find:

$$\Theta(\lambda, x_0) = \frac{a(1 - a^2)(1 - x_0^2)^2}{(1 - x_0^2 + a^2 x_0^2)^{3/2} + a(1 - a^2)(1 - x_0^2)^2 \lambda} \quad (3.12)$$

In Fig. 8 we have plotted this as a function of λ for $x_0 = 0$, at which $\Theta = \frac{a(1 - a^2)}{1 + a(1 - a^2)\lambda}$.

For $a > 1$, we expect the congruence to develop a caustic where the expansion diverges. This occurs when infinitesimally nearby geodesics intersect each other. Eq.

(3.12) shows that this can only occur for $a > 1$, where the second term in the denominator is negative for positive λ . In this case $\Theta(\lambda) \rightarrow -\infty$ at a finite value of $\lambda = \lambda_c$,

$$\lambda_c = \frac{(1 - x_0^2 + a^2 x_0^2)^{3/2}}{a(a^2 - 1)(1 - x_0^2)} \quad (3.13)$$

for any x_0 . The spacetime coordinates for the points along the congruence where this happens are given by

$$x_c = (1 - a^2)x_0^3, \quad t_c = \frac{(1 - x_0^2 + a^2 x_0^2)^{3/2}}{a}, \quad z_c = \frac{a^2 - 1}{a}(1 - x_0^2)^{3/2} \quad (3.14)$$

Viewed as a pair of parametric curves parametrized by x_0 which starts at $x_0 = 0$ and ends at $x_0 = \pm 1$, the caustic seams are null curves, starting at the intersection point (3.6) and ending on the boundary at $z_c = 0$, $x_c = \pm(1 - a^2)$, and $t_c = a^2$. Note that this is a finite distance on the boundary.

The divergence $\Theta \rightarrow -\infty$ signifies the presence of conjugate points, but their geometric meaning is a bit obscure in our discussion so far. The reason is as follows: as we see in Fig. 6 and can check explicitly, we generically have caustics in the neighbourhood of $x_0 \simeq 0$, but more generally encounter cross-over points from the intersection geodesics symmetrically placed about $x_0 = 0$. The expansion is *finite* along the cross-over seam (3.5) for $x_0 \neq 0$. This can be understood by realizing that the expansion is a local property of the nearby geodesics which doesn't know about any other piece of the congruence. So nothing special ought to happen at the cross-over points which are non-local in the congruence, and indeed these are not conjugate points.

The clue as to the geometric meaning of $\Theta \rightarrow -\infty$ comes from plotting this locus on the surface of the null congruence (continued through the cross-over seam). This is presented in Fig. 9 by the thick red curves. We see that the surface intersects itself at the cross-over seam, beyond which the constant- λ wavefronts form closed loops. On the sharp flank, these wavefronts turn around and locally become null; this is precisely where $A(\lambda, x_0)$ vanishes and therefore $\Theta \rightarrow -\infty$.

3.4 Summary

The upshot of our calculations can be summarized as follows. Consider the null geodesic congruence emanating from a codimension-two spacelike surfaces $\mathcal{F}_{\mathcal{A}} \subset \mathcal{M}$ anchored on the boundary of a region \mathcal{A} with $\partial\mathcal{A} = \mathcal{F}_{\mathcal{A}} \cap \mathcal{B}$.

- If $\mathcal{F}_{\mathcal{A}} \subset \mathcal{W}_{\mathcal{C}}[\mathcal{A}]$ then the congruence terminates inside $D[\mathcal{A}]$ along a spacelike boundary codimension-one surface.




Fig. 9: Surface generated by the null normal congruence, along with the locus of points on this surface where the expansion diverges, indicated by the thick red curves. The cyan contours represent the geodesic generators, while the blue contours are the constant- λ wavefronts (we cut off the surface at $|x_0| < 1$ for convenience).

- If $\mathcal{F}_{\mathcal{A}}$ lies on the boundary of the causal wedge $\mathcal{W}_{\mathcal{C}}[\mathcal{A}]$ then the congruence intersects the boundary on the null surface $\partial D[\mathcal{A}]$.
- If $\mathcal{F}_{\mathcal{A}} \subset \tilde{\mathcal{S}}[D[\mathcal{A}]]$ then the congruence finds itself terminated by a seam of cross-over points (and if continued further, would encounter caustic points prior to reaching the AdS boundary). The seam itself however reaches out to the boundary and ends on the future tip¹⁸ of $\partial D[\mathcal{A}]$.

This gives a clear picture of the causal domains for regions bounded by curves inside and outside of $\mathcal{W}_{\mathcal{C}}[\mathcal{A}]$. As we will see in our explicit proof, the extremal surface will in general lie outside of $\mathcal{W}_{\mathcal{C}}[\mathcal{A}]$; in special cases it can at best lie on the boundary, but never in the interior, of the causal wedge.

4 Theorem and proof

We now get to the main part of the paper where we prove that the extremal surface $\mathcal{E}_{\mathcal{A}}$ satisfies the causality requirements discussed in §2.3. Our main goal will be to establish

¹⁸ In higher-dimensional setting, $D[\mathcal{A}]$ itself may terminate in a crossover seam rather than a single point, which occurs when the null generators of $\partial D[\mathcal{A}]$ on the boundary themselves cross over.

the causal relations quoted there in (2.11). These will establish for us the consistency of the HRT proposal for computing holographic entanglement entropy.

In §4.1, we remind the reader of the holographic set-up and of our assumptions. In §4.2, we study null geodesic congruences in the bulk and their intersections with the boundary. In particular, since a geodesic that reaches the boundary travels an infinite affine parameter, a non-expanding congruence that reaches the boundary without hitting a caustic must have vanishing shear, and therefore must intersect the boundary at a single point. This allows us to show, using the null energy condition, that the intersection with the boundary of the causal future of an extremal bulk surface equals the causal future of its intersection with the boundary. As a warm-up, we prove a version of the Gao-Wald theorem [14]. Finally, in §4.3, we carefully define what we mean by a region and by the spacelike homology condition. We prove that a region \mathcal{A} implies a natural decomposition of the spacetime into four regions $D[\mathcal{A}]$, $D[\mathcal{A}^c]$, and $J^\pm[\partial\mathcal{A}]$. Then, given the spacelike homology condition, and using the results of §4.2, we establish the compatibility of the boundary and bulk decompositions, (2.11), and prove that the extremal surface is a wedge observable.

4.1 Holographic setup

In this subsection we will describe our holographic setup and assumptions.¹⁹

Let (\mathcal{M}, g_{ab}) be a connected spacetime, of dimension greater than or equal to 3, that can be embedded in a spacetime $(\bar{\mathcal{M}}, \tilde{g}_{ab})$, such that the boundary \mathcal{B} of \mathcal{M} in $\bar{\mathcal{M}}$ is a smooth timelike hypersurface in $\bar{\mathcal{M}}$, and such that $\tilde{g}_{ab} = \Omega^2 g_{ab}$, where Ω is a smooth function on $\bar{\mathcal{M}}$ that vanishes on \mathcal{B} . (We do not assume that \mathcal{B} is connected.) We define $\tilde{\mathcal{M}} := \mathcal{M} \cup \mathcal{B}$. On $\tilde{\mathcal{M}}$ we have a causal structure induced from \tilde{g}_{ab} , which in \mathcal{M} agrees with that induced from g_{ab} . We make the following assumptions:

- (i) (\mathcal{M}, g_{ab}) obeys the null energy condition.
- (ii) $\tilde{\mathcal{M}}$ is globally hyperbolic.
- (iii) Every null geodesic in $(\mathcal{B}, \tilde{g}_{ab})$ is a geodesic in $(\tilde{\mathcal{M}}, \tilde{g}_{ab})$.²⁰

¹⁹ We largely follow the setup and assumptions of section 3 of [14], with two exceptions: we remove the null generic condition and we add the condition that the boundary is totally geodesic for null geodesics (assumption (iii) below).

²⁰ Assumption (iii) is equivalent to the following property of the extrinsic curvature K_{ab} of \mathcal{B} in $\tilde{\mathcal{M}}$: for any point $p \in \mathcal{B}$ and any null vector k^a in the tangent space to \mathcal{B} at p , $K_{ab}k^a k^b = 0$. That it holds for an asymptotically AdS spacetime can be seen by working in Fefferman-Graham coordinates. If we set $\Omega = 1/z$, where z is the standard radial coordinate, then $K_{ab} = 0$ (so all geodesics in \mathcal{B} are geodesics in $\tilde{\mathcal{M}}$, i.e. \mathcal{B} is totally geodesic). The property $K_{ab} = 0$ is not preserved by Weyl transformations, and

We begin by showing that \mathcal{B} is globally hyperbolic. We omit the proofs, which are very simple, cf., [40]. (For brevity, we will only indicate one time direction for each statement below, but the time-reversed statements are clearly equally valid.)

Lemma 1 *For any set $\Upsilon \subset \tilde{\mathcal{M}}$, $\tilde{D}^+[\Upsilon] \cap \mathcal{B} \subset D[\Upsilon \cap \mathcal{B}]$.*

Lemma 2 *If $\tilde{\Sigma} \subset \tilde{\mathcal{M}}$ is closed and acausal, then $\tilde{\Sigma} \cap \mathcal{B}$ is closed and acausal in \mathcal{B} .*

Corollary 3 *If $\tilde{\Sigma}$ is a Cauchy slice²¹ for $\tilde{\mathcal{M}}$, then $\tilde{\Sigma} \cap \mathcal{B}$ is a Cauchy slice for \mathcal{B} .*

Corollary 4 *\mathcal{B} is globally hyperbolic.*

4.2 Congruences of null geodesics

In this subsection, we will study null geodesics in $\tilde{\mathcal{M}}$. Assumption (iii) has the following useful implication:

Lemma 5 *Any null geodesic in $\tilde{\mathcal{M}}$ either (1) lies entirely in \mathcal{B} , or (2) does not intersect \mathcal{B} except possibly at its endpoints, where it is not tangent to \mathcal{B} .*

Proof: Given a point p in \mathcal{B} and a non-zero null vector in the tangent space to \mathcal{B} at p , there exists a null geodesic in \mathcal{B} passing through p with that tangent vector. By assumption (iii), it is a geodesic in $\tilde{\mathcal{M}}$, and by the uniqueness of geodesics it is the only one. Therefore no null geodesic passing through \mathcal{M} can intersect \mathcal{B} tangentially. Finally, since \mathcal{B} is the boundary of $\tilde{\mathcal{M}}$ and is smooth, any smooth curve that intersects \mathcal{B} at some point without ending there must be tangent to it. \square

Now we constrain the behavior of congruences of null geodesics that pass through \mathcal{M} , using the fact that the metric g_{ab} obeys the null energy condition and the fact that a geodesic that reaches \mathcal{B} travels an infinite affine parameter.

Lemma 6 *Consider a codimension-one congruence of future-directed null geodesics in $\tilde{\mathcal{M}}$, each of which lies entirely in \mathcal{M} except possibly at its endpoints. Suppose that the part of the congruence in \mathcal{M} has the following properties: (1) its expansion with respect to the metric g_{ab} is nowhere positive; (2) at each point, every deviation vector is spacelike and orthogonal to the tangent vector. Then the congruence intersects \mathcal{B} on a set of isolated points.*

so does not hold for a general choice of Ω , but the weaker condition $K_{ab}k^ak^b = 0$ does (as can be seen either from a direct calculation or from the fact that the set of null geodesics is invariant under Weyl transformations).

²¹ We remind the reader that, as explained in footnote 4, throughout this paper we require all Cauchy slices to be acausal, not just achronal.

Proof: We begin by working in the metric g_{ab} . Since the deviation vectors are everywhere spacelike, the expansion Θ is finite everywhere. On any geodesic that reaches \mathcal{B} , the affine parameter goes to infinity, so, by the null energy condition, Θ is nowhere negative, and therefore vanishes everywhere. Again using the null energy condition, the shear therefore vanishes everywhere also. Therefore, for any one-parameter family of geodesics that reach \mathcal{B} , the norm of the deviation vector X^a is a positive constant along each geodesic.

We now return to $\tilde{\mathcal{M}}$, and switch to the metric \tilde{g}_{ab} . On \mathcal{B} , X^a has vanishing norm; being also orthogonal to the geodesic's tangent vector T^a , it is proportional to T^a (since orthogonal null vectors are proportional). Without loss of generality, we choose the affine parameter λ on each geodesic so that it intersects \mathcal{B} at $\lambda = 0$; hence, at $\lambda = 0$, X^a is tangent to \mathcal{B} . However, by lemma 5, T^a is not tangent to \mathcal{B} . So $X^a = 0$. Since this holds for every one-parameter family of geodesics, every connected set of geodesics that reach \mathcal{B} intersects it at a point. \square

As a warm-up for our main theorem of this subsection, we will now use lemma 6 to prove a version of the Gao-Wald theorem [14] and a version of the topological censorship theorem [41].

Theorem 7 *For any point $p \in \mathcal{B}$, $\tilde{J}^+(p) \cap \mathcal{B} = J^+(p)$.*

Proof: Clearly $J^+(p) \subset \tilde{J}^+(p) \cap \mathcal{B}$. Let t be a global time function on $\tilde{\mathcal{M}}$. Then if $t(q) < t(p)$ we have $q \notin \tilde{J}^+(p)$. Therefore, each connected component of \mathcal{B} contains some points not in $\tilde{J}^+(p)$. Therefore, if $\tilde{J}^+(p) \cap \mathcal{B} \neq J^+(p)$, then $\partial\tilde{J}^+(p) \cap \mathcal{B}$ includes a hypersurface \mathcal{S} in \mathcal{B} that is not in $J^+(p)$. We will now show that \mathcal{S} cannot exist.

$\partial\tilde{J}^+(p)$ consists of future-directed null geodesics starting at p on which, except at the endpoints, every deviation vector is spacelike and orthogonal to the tangent vector. By lemma 5, each such geodesic either lies entirely in \mathcal{B} or lies entirely in \mathcal{M} except at its endpoints. In particular, the points in \mathcal{S} must lie on geodesics that are entirely in \mathcal{M} except at their endpoints. We thus consider the congruence of geodesics in \mathcal{M} starting at p . Reversing its direction, every geodesic in this congruence reaches \mathcal{B} (at p), so the expansion is nowhere negative. Therefore, in the forward direction, its expansion is nowhere positive. Thus the conditions of lemma 6 apply. Hence \mathcal{S} consists of isolated points, contradicting the fact that it is a hypersurface in \mathcal{B} . \square

Corollary 8 *If $\mathcal{B}_1, \mathcal{B}_2$ are distinct connected components of \mathcal{B} , then $\tilde{J}^+(\mathcal{B}_1) \cap \mathcal{B}_2 = \emptyset$.*

Corollary 8 rules out traversable wormholes through the bulk connecting different boundary components, and is thus closely related to topological censorship. (A simple argument establishing this can be found in [42].)

Our goal for the rest of this subsection is generalize Theorem 7 to codimension-two surfaces that are extremal with respect to g_{ab} . First, we need two lemmas:

Lemma 9 *Let \mathcal{E} be a compact codimension-two submanifold-with-boundary of $\tilde{\mathcal{M}}$, with boundary \mathcal{N} . Then every point $p \in \partial\tilde{J}^+[\mathcal{E}]$ is on a future-directed null geodesic lying entirely in $\partial\tilde{J}^+[\mathcal{E}]$ that either (1) starts orthogonally from \mathcal{E} and has no point conjugate to \mathcal{E} between \mathcal{E} and p , or (2) starts orthogonally from \mathcal{N} , moving away from \mathcal{E} (i.e. $U_a T^a > 0$, where T^a is the tangent vector to the geodesic at its starting point, and U^a is a vector at the same point that is tangent to \mathcal{E} , normal to \mathcal{N} , and outward-directed from \mathcal{E}).*

Proof: This is a generalization of theorem 9.3.11 in [25]. Every $p \in \partial\tilde{J}^+[\mathcal{E}]$ lies on a null geodesic starting from \mathcal{E} . If neither condition (1) nor (2) is met, then it can be deformed to a timelike curve and therefore $p \in \tilde{I}^+[\mathcal{E}]$. \square

Lemma 10 *Let \mathcal{E} be a spacelike submanifold-with-boundary of $\tilde{\mathcal{M}}$ whose restriction to \mathcal{M} is extremal with respect to the metric g_{ab} . Then \mathcal{E} intersects \mathcal{B} orthogonally, i.e., every normal vector to \mathcal{E} is tangent to \mathcal{B} .*

Proof: A short calculation shows that, in \mathcal{M} , the mean curvature \tilde{K}^a of \mathcal{E} with respect to \tilde{g}_{ab} is related to that with respect to g_{ab} , K^a , as follows:

$$\tilde{K}^a = \Omega^{-2} K^a + \dim(\mathcal{E}) \tilde{Q}^{ab} \partial_b \ln \Omega, \quad (4.1)$$

where $\tilde{Q}^{ab} := Q^a_c \tilde{g}^{bc}$ and Q^a_c is the projector normal to \mathcal{E} . Since \mathcal{E} is extremal, $K^a = 0$. So

$$\tilde{K}^2 = \dim(\mathcal{E})^2 \tilde{Q}^{ab} \partial_a \ln \Omega \partial_b \ln \Omega. \quad (4.2)$$

Since \mathcal{E} is smooth, \tilde{K}^2 remains finite on \mathcal{B} , where $\ln \Omega \rightarrow -\infty$. This requires that every normal vector to \mathcal{E} be tangent to \mathcal{B} . \square

Theorem 11 *Let \mathcal{E} be a compact smooth spacelike codimension-two submanifold-with-boundary in $\tilde{\mathcal{M}}$, whose only boundary is where it intersects \mathcal{B} , and whose restriction to \mathcal{M} is extremal with respect to the metric g_{ab} . Then $\tilde{J}^+[\mathcal{E}] \cap \mathcal{B} = J^+[\mathcal{E} \cap \mathcal{B}]$.*

Proof: The proof is largely a repetition of that of Theorem 7. Clearly $J^+[\mathcal{E} \cap \mathcal{B}] \subset \tilde{J}^+[\mathcal{E}] \cap \mathcal{B}$. Let t be a global time function on $\tilde{\mathcal{M}}$. Since \mathcal{E} is compact, it has a minimum time t_{\min} . Clearly if for some point $q \in \mathcal{B}$, $t(q) < t_{\min}$, then $q \notin \tilde{J}^+[\mathcal{E}]$. Therefore, each connected component of \mathcal{B} contains some points not in $\tilde{J}^+[\mathcal{E}]$. Therefore, if $\tilde{J}^+[\mathcal{E}] \cap \mathcal{B} \neq J^+[\mathcal{E} \cap \mathcal{B}]$, then $\partial\tilde{J}^+(m) \cap \mathcal{B}$ includes a hypersurface Σ in \mathcal{B} that is not in $\tilde{J}^+[\mathcal{E} \cap \mathcal{B}]$. We will now show that \mathcal{S} cannot exist.

By lemma 10, \mathcal{E} intersects \mathcal{B} orthogonally. Therefore, in lemma 9, the second type of null geodesic in $\partial\tilde{J}^+[\mathcal{E}]$ does not exist. The first type of geodesic forms a codimension-two congruence starting orthogonally from \mathcal{E} on which, except possibly at the endpoints, every deviation vector is spacelike and orthogonal to the tangent vector. By lemma 5, each such geodesic either lies entirely in \mathcal{B} or lies entirely in \mathcal{M} except at its endpoints. In particular, the points in \mathcal{S} must lie on geodesics that are entirely in \mathcal{M} except where they end. We thus consider the congruence of geodesics in \mathcal{M} starting orthogonally from $\mathcal{E} \cap \mathcal{M}$. Since $\mathcal{E} \cap \mathcal{M}$ is extremal, its expansion (with respect to g_{ab}) is initially zero. By the null energy condition, its expansion is nowhere positive. Thus the conditions of lemma 6 apply. Hence \mathcal{S} consists of isolated points, contradicting the fact that it is a hypersurface in \mathcal{B} . \square

Note that theorem 7 is a special case of theorem 11, in which we take \mathcal{E} to be a small (in the metric \tilde{g}_{ab}) hemisphere centered on p and take the limit in which its radius goes to 0.

4.3 Spatial regions and causal decompositions

Let Σ be a Cauchy slice of \mathcal{B} . Given a codimension-zero submanifold of Σ , let \mathcal{A} be its interior, $\partial\mathcal{A}$ its boundary, and \mathcal{A}^c its complement; these three sets do not overlap and cover Σ . They naturally induce a causal decomposition of the spacetime \mathcal{B} into four nonoverlapping regions (except that $J^\pm[\partial\mathcal{A}]$ both include $\partial\mathcal{A}$):

Theorem 12

$$D[\mathcal{A}] \cup D[\mathcal{A}^c] \cup J^+[\partial\mathcal{A}] \cup J^-[\partial\mathcal{A}] = \mathcal{B} \quad (4.3)$$

$$D[\mathcal{A}] \cap D[\mathcal{A}^c] = D[\mathcal{A}] \cap J^\pm[\partial\mathcal{A}] = D[\mathcal{A}^c] \cap J^\pm[\partial\mathcal{A}] = \emptyset \quad (4.4)$$

$$J^+[\partial\mathcal{A}] \cap J^-[\partial\mathcal{A}] = \partial\mathcal{A}. \quad (4.5)$$

Proof: (4.4) and (4.5) are obvious from the definitions.

We now prove (4.3). Suppose a point $p \in J^+[\Sigma]$ is not in any of the four regions. Each inextendible causal curve through p intersects Σ exactly once, but not in $\partial\mathcal{A}$ (else $p \in J^+[\partial\mathcal{A}]$). Nor can all such curves intersect it in \mathcal{A} (else $p \in D[\mathcal{A}]$) or \mathcal{A}^c (else $p \in D[\mathcal{A}^c]$). So some must intersect Σ in \mathcal{A} and others in \mathcal{A}^c . Let λ_1 be in the first set and λ_2 in the second. Join λ_1 and λ_2 at p to make a continuous curve λ from \mathcal{A} to \mathcal{A}^c . Now, in any globally hyperbolic spacetime there exists a global timelike vector field; its integral curves can be used to construct a continuous map f from $J^+(\Sigma)$ to Σ . $f(\lambda)$ is a continuous curve in Σ from \mathcal{A} to \mathcal{A}^c . There therefore exists a point $q \in \lambda$

such that $f(q) \in \partial\mathcal{A}$, and therefore $q \in I^+[\partial\mathcal{A}]$. Since $p \in J^+(q)$, $p \in J^+[\partial\mathcal{A}]$, which is a contradiction. \square

Now let $\mathcal{E}_{\mathcal{A}}$ be a surface in $\tilde{\mathcal{M}}$ that satisfies the conditions of theorem 11 and is spacelike-homologous to \mathcal{A} . The precise meaning of the latter condition is as follows: There exists a Cauchy slice $\tilde{\Sigma}$ for $\tilde{\mathcal{M}}$ such that $\tilde{\Sigma} \cap \mathcal{B} = \Sigma$, containing a codimension-zero submanifold with boundary $\mathcal{A} \cup \mathcal{E}_{\mathcal{A}}$; we call its interior $\mathcal{R}_{\mathcal{A}}$. Since $\tilde{\Sigma}$ is itself a manifold-with-boundary (namely $\tilde{\Sigma} \cap \mathcal{B}$), one has to be careful about the definitions of “interior” and “boundary” for a submanifold. We mean “interior” in the sense of point-set topology; thus $\mathcal{R}_{\mathcal{A}}$ includes \mathcal{A} but not $\mathcal{E}_{\mathcal{A}}$. The “boundary” can be either in the sense of “submanifold-with-boundary” (which is what we call $\partial\mathcal{R}_{\mathcal{A}}$), or in the sense of point-set topology. In the latter sense, the boundary is just $\mathcal{E}_{\mathcal{A}}$.²² As with \mathcal{A} , we define $\mathcal{R}_{\mathcal{A}}^c := \tilde{\Sigma} \setminus (\mathcal{R}_{\mathcal{A}} \cup \mathcal{E}_{\mathcal{A}})$. To summarize, in parallel to the decomposition of Σ into \mathcal{A} , \mathcal{A}^c , and $\partial\mathcal{A}$, we have a decomposition of $\tilde{\Sigma}$ into $\mathcal{R}_{\mathcal{A}}$, $\mathcal{R}_{\mathcal{A}}^c$, and $\mathcal{E}_{\mathcal{A}}$. Furthermore, $\mathcal{R}_{\mathcal{A}} \cap \mathcal{B} = \mathcal{A}$, $\mathcal{R}_{\mathcal{A}}^c \cap \mathcal{B} = \mathcal{A}^c$, and $\mathcal{E}_{\mathcal{A}} \cap \mathcal{B} = \partial\mathcal{A}$.

We can now apply theorem 12 to obtain a decomposition of $\tilde{\mathcal{M}}$ into the four spacetime regions $D[\mathcal{R}_{\mathcal{A}}]$, $D[\mathcal{R}_{\mathcal{A}}^c]$, $J^\pm[\mathcal{E}_{\mathcal{A}}]$. The central result of this section is that this decomposition reduces on the boundary precisely to its decomposition into $D[\mathcal{A}]$, $D[\mathcal{A}^c]$, and $J^\pm[\partial\mathcal{A}]$:

Theorem 13

$$\tilde{D}[\mathcal{R}_{\mathcal{A}}] \cap \mathcal{B} = D[\mathcal{A}] \tag{4.6a}$$

$$\tilde{D}[\mathcal{R}_{\mathcal{A}}^c] \cap \mathcal{B} = D[\mathcal{A}^c] \tag{4.6b}$$

$$\tilde{J}^\pm[\mathcal{E}_{\mathcal{A}}] \cap \mathcal{B} = J^\pm[\partial\mathcal{A}] \tag{4.6c}$$

Proof: Equation (4.6c) is Theorem 11 (and its time reverse). Using Theorem 12 both in \mathcal{B} and in $\tilde{\mathcal{M}}$ to take the complement of both sides, we have

$$\left(\tilde{D}[\mathcal{R}_{\mathcal{A}}] \cap \mathcal{B}\right) \cup \left(\tilde{D}[\mathcal{R}_{\mathcal{A}}^c] \cap \mathcal{B}\right) = D[\mathcal{A}] \cup D[\mathcal{A}^c]. \tag{4.7}$$

Lemma 1 then implies (4.6a), (4.6b). \square

Theorem 13 immediately implies that $\mathcal{E}_{\mathcal{A}}$ is outside of causal contact with $D[\mathcal{A}]$ and $D[\mathcal{A}^c]$, as required by field-theory causality.

The spacelike-homology condition raises the following practical question: Given a codimension-one submanifold of $\tilde{\mathcal{M}}$ with boundary $\mathcal{A} \cup \mathcal{E}_{\mathcal{A}}$, under what circumstances

²² The point-set-topology boundary can be shown to equal the “edge” of the submanifold, in the sense used in the general-relativity literature (see e.g. [25]).

is it contained in a Cauchy slice? Obviously, it must be acausal. However, this is not sufficient; for example, a spacelike hypersurface in Minkowski space that approaches null infinity is not contained in a Cauchy slice. The following lemma, which will also be needed in theorem 15, shows that compactness is a sufficient additional condition. (This lemma applies in any globally hyperbolic spacetime.)

Lemma 14 *If R is a compact acausal set, then there exists a Cauchy slice containing it.*

Proof: Let $t \in \mathbb{R}$ be a global time function, and define $t_{\max} := \max_R(t)$, $t_{\min} := \min_R(t)$ (these exist since R is compact). Define $\Upsilon := \{p : t > t_{\max}\}$ and $\Upsilon' := \Upsilon \cup I^+[R]$. Define

$$\Sigma := \partial\Upsilon' = (\partial\Upsilon \setminus I^+[R]) \cup (\partial I^+[R] \setminus \Upsilon). \quad (4.8)$$

$\partial I^+[R]$ contains R , and $\Upsilon \cap R = \emptyset$, so $R \subset \partial I^+[R] \setminus \Upsilon \subset \Sigma$. Next we show that Σ is achronal. The maximum value of t on Σ is t_{\max} , so there can be no future-directed timelike curve from $\partial\Upsilon$ to Σ . Further $\partial I^+[R]$ is itself achronal. Finally, if there is a future-directed timelike curve from $p \in \partial I^+[R]$ to $q \in \partial\Upsilon$, then $q \in I^+[R]$ and hence $q \notin \Sigma$. So Σ is achronal.

Next, we show that every inextendible future-directed timelike curve intersects Σ . On such a curve, t increases monotonically and continuously from $-\infty$ to $+\infty$. For $t \leq t_{\min}$, the curve is not in Υ' ; for $t > t_{\max}$, it is. Therefore for some value of t it intersects Σ .

While Σ is achronal, it is not quite a Cauchy slice (in the sense used in this paper) because it is not acausal. However, since R is acausal, Σ can be deformed outside of R to be acausal. \square

Theorem 15 *Let Σ' be a Cauchy slice for \mathcal{B} and $\mathcal{A}' \subset \Sigma'$ a region such that $\mathcal{A}' \cup \partial\mathcal{A}'$ is compact and $D[\mathcal{A}'] = D[\mathcal{A}]$. Then \mathcal{A}' is spacelike-homologous to $\mathcal{E}_{\mathcal{A}}$.*

Proof: Since $\mathcal{E}_{\mathcal{A}}$ and $\mathcal{A}' \cup \partial\mathcal{A}'$ are both compact, $\mathcal{E}_{\mathcal{A}} \cup \mathcal{A}'$ is compact as well. (Recall that $\partial\mathcal{A}' = \partial\mathcal{A} \subset \mathcal{E}_{\mathcal{A}}$.) $\mathcal{E}_{\mathcal{A}}$ and \mathcal{A}' are acausal, since each sits on a Cauchy slice. Furthermore, by theorems 12 and 11, there are no causal curves connecting them; hence $\mathcal{E}_{\mathcal{A}} \cup \mathcal{A}'$ is acausal. Therefore, by theorem 14, there is a Cauchy slice $\tilde{\Sigma}'$ containing both $\mathcal{E}_{\mathcal{A}}$ and \mathcal{A}' .

Choosing a global timelike vector field on $\tilde{\mathcal{M}}$, its integral curves define a diffeomorphism $f : \tilde{\Sigma} \rightarrow \tilde{\Sigma}'$. Let $\mathcal{R}'_{\mathcal{A}} := f(\mathcal{R}_{\mathcal{A}})$. Since $\mathcal{E}_{\mathcal{A}}$ is contained in both Σ and Σ' , $f(\mathcal{E}_{\mathcal{A}}) = \mathcal{E}_{\mathcal{A}}$. Since every timelike curve in $D[\mathcal{A}]$ intersects Σ in \mathcal{A} and Σ' in \mathcal{A}' ,

$f(\mathcal{A}) = \mathcal{A}'$. So $\mathcal{R}'_{\mathcal{A}} := f(\mathcal{R}_{\mathcal{A}})$ is a region in $\tilde{\Sigma}'$ with $\partial\mathcal{R}'_{\mathcal{A}} = \mathcal{A}' \cup \mathcal{E}_{\mathcal{A}}$. (Strictly speaking, we also need to define a new Cauchy slice for \mathcal{B} , $\Sigma'' := \tilde{\Sigma}' \cap \mathcal{B}$, and to consider \mathcal{A}' to be a region in Σ'' , since the equality $\Sigma'' = \tilde{\Sigma}' \cap \mathcal{B}$ is part of the definition of the spacelike homology condition.) \square

Theorem 15 shows that the HRT formula gives the same value for the entanglement entropy of \mathcal{A} and \mathcal{A}' , as required by field-theory causality.

5 Discussion

The main result of this paper, Theorem 13, shows that the HRT prescription for computing holographic entanglement entropy [32] is consistent with the requirements of field theory causality. As we have explained with various simple examples and gedanken experiments in §2.4, the result was in no way *a priori* obvious, since there are several marginal cases where arbitrarily small deformation of the bulk extremal surface would place it in causal future of a boundary deformation which however cannot affect the entanglement entropy. With the primary result at hand, we now take stock of the various physical consequences it implies for holographic field theories.

Causality constraints on holography: Let us start by asking what we can learn about holography from causality considerations. Recall that we proved our result for extremal surfaces in the context of two-derivative theories of gravity satisfying the null energy condition. This was crucial for us to be able to use the Raychaudhuri equation in order to ascertain properties of null geodesic congruences. Thus the domain of validity of our statements was strong coupling in a planar (large- N) field theory. This translates to demanding a macroscopic spacetime with $\ell_s \ll \ell_{\text{AdS}}$ in a perturbative string ($g_s \ll 1$) regime. Lets see what happens as we move away from this corner of moduli space.

Firstly, consider classical stringy corrections which we can encapsulate in an effective higher-derivative theory of gravity. In such a theory, as long as higher-derivative operators are suppressed by powers of ℓ_s , our conclusions will hold, since the dominant effect will come from the leading two-derivative Einstein-Hilbert term in the bulk. When the higher-derivative operators are unsuppressed we have little to say for two reasons: (a) the holographic entanglement prescription so far is only given for static situations (or with time reversal symmetry) [43, 44] and (b) even assuming the covariant generalizations, one is stymied by the absence of clean statements regarding dynamics of null geodesic congruences (even for example in Lovelock theories).²³ One

²³ The family of $f(R)$ theories can be brought to heel, since here we can map the theory to Einstein-Hilbert via a suitable Weyl transformation. Causality constraints can be discerned here so long as the Weyl transformation (which is non-linear in the curvature) is well-behaved.

could, however, use the causality constraint to rule out certain higher-derivative theories from having unitary relativistic QFT duals (see e.g. [45]); this is similar in spirit to the recent discussions on causality constraints on the three-graviton vertex [46].

Turning next to $1/N$, or bulk quantum corrections, while we have less control in general, we can make some observations about the leading $1/N$ correction which has been proposed to be given by the entanglement of bulk perturbative quantum fields across \mathcal{E}_A [47]. Since the bulk theory itself is causal, it follows that entanglement across the extremal surface satisfies the desired causality conditions.

Does causality prove the HRT conjecture?: One intriguing possibility given, the importance of the causality, is whether we can use it to constrain the location of the extremal surface in the bulk, and thus prove the HRT conjecture.²⁴ Unfortunately, causality alone is not strong enough to pin down the location of the extremal surface. What we can say is that the extremal surface \mathcal{E}_A has to lie inside the causal shadow $\mathcal{Q}_{\partial\mathcal{A}}$. In a generic asymptotic AdS spacetime, for a generic region \mathcal{A} , the casual shadow is a codimension-zero volume of the bulk spacetime \mathcal{M} . It is only in some very special cases that we zero in on a single bulk codimension-two surface uniquely (e.g., spherical regions in pure AdS or in the eternal Schwarzschild-AdS black hole).²⁵

Causality constraints on other CFT observables: Our discussion has exclusively focused on the causality properties of a particular non-local quantity in the field theory, namely the entanglement entropy. However, causality places restrictions on other physical observables we can consider on the boundary as well. For instance, correlation functions of (time-ordered) local operators, Wilson loop expectation values, etc., should all obey appropriate constraints which we can infer from basic principles. Indeed, this can be shown to be the case, for example, for correlation functions, by considering the fact that the bulk computation involves solving a suitable boundary initial value problem for fields in the bulk, which can be checked to manifestly satisfy causality.

However, this is less clear when we approximate, say, two point functions of heavy local operators using the geodesic approximation [48]. Similar issues arise for the semi-classical computation of Wilson loop expectation values [49, 50] using the string worldsheet area. In these cases, one generically encounters some tension between the use of extremal surfaces—geodesics, two-dimensional worldsheets, etc.—for the bulk computation, and field theory expectations regarding causality (cf., [51] for an earlier

²⁴ We thank Vladimir Rosenhaus for inspiring us to think through this possibility.

²⁵ The examples are all cases where, by a suitable choice of conformal frame, the extremal surface can be mapped onto the bifurcation surface of a static black hole. The black funnel and droplet solutions (see [30] for a review) provide nontrivial examples, cf., [23].

discussion of this issue). Indeed, it appears that codimension-two extremal surfaces are special in this regard, for we can rely on the boundary of the entanglement wedge being generated by a codimension-one null congruence, and thus apply the Raychaudhuri equation. Understanding the proper application of the WKB approximation for other observables is an interesting question; we hope to report upon in the near future [52].

Entanglement wedges: One of the key constructs in our presentation, naturally associated with a given boundary region \mathcal{A} , has been the entanglement wedge $\mathcal{W}_{\mathcal{E}}[\mathcal{A}]$. This is the domain of dependence of the homology surface $\mathcal{R}_{\mathcal{A}}$ (recall that $\mathcal{R}_{\mathcal{A}}$ forms a part of a Cauchy surface which interpolates between \mathcal{A} and $\mathcal{E}_{\mathcal{A}}$). Equivalently, it comprises the set of spacelike-separated points from $\mathcal{E}_{\mathcal{A}}$ which is connected to \mathcal{A} , one of the four regions in the natural decomposition of the bulk spacetime.

Given \mathcal{A} , one might ask how unique this decomposition is. Since $\mathcal{W}_{\mathcal{E}}[\mathcal{A}]$ is a causally-defined set, its specification only requires the specification of the (oriented) extremal surface $\mathcal{E}_{\mathcal{A}}$ (possibly consisting of multiple components when so required by the homology constraint). The prescription for constructing the null boundary of $\mathcal{W}_{\mathcal{E}}[\mathcal{A}]$ is unambiguous: simply to follow all null normals (emanating from $\mathcal{E}_{\mathcal{A}}$ in the requisite direction, towards $D[\mathcal{A}]$) until they encounter another generator (i.e. a crossover seam) or a caustic. However, there is a possibility that the extremal surface itself is not uniquely determined from \mathcal{A} . This happens when multiple (sets of) extremal surfaces satisfy (2.4) but have the same area. Since entanglement entropy itself cares only about the area, the HRT (as well as RT and maximin) prescription is to take *any* of these. However, which we take *does* matter for the entanglement wedge. We propose that, just as for the extremal surfaces, in such cases we may have *multiple* entanglement wedges $\mathcal{W}_{\mathcal{E}}[\mathcal{A}]$ associated to the same boundary region \mathcal{A} .

The most “obvious” class of examples where this can happen is the case of \mathcal{A} consisting of multiple regions or in higher dimensions where the entangling surface $\partial\mathcal{A}$ consists of multiple disjoint components. As we vary the parameters describing the configuration, the extremal surfaces involved typically exchange dominance, so at some point their areas must agree. Applying continuity from both sides, at the transition point, both entanglement wedges should be naturally associated with \mathcal{A} . However, in complicated states, there can actually be multiple extremal surfaces even for when \mathcal{A} and $\partial\mathcal{A}$ are both connected. In such cases, we could have candidate entanglement wedges which are proper subsets of (rather than merely overlapping with) other candidate entanglement wedges.

It is also interesting to note that the decomposition of the bulk into four spacetime regions causally defined from $\mathcal{E}_{\mathcal{A}}$ need not coincide with the bulk decomposition defined from $\mathcal{E}_{\mathcal{A}^c}$, despite there being a unique boundary decomposition defined from $\partial\mathcal{A}$. For




Fig. 10: Sketch of Penrose diagram for a symmetric Vaidya-Schwarzschild-AdS geometry obtained by imploding null shells to the past and future from both boundaries now displaying the entanglement wedges and the causal shadow region, with \mathcal{A} being a full Cauchy surface for CFT_R .

pure states, where the homology constraint trivializes and we have $\mathcal{E}_A = \mathcal{E}_{A^c}$, we can write the bulk decomposition equivalently with respect to both \mathcal{A} and \mathcal{A}^c ,

$$\mathcal{M} = \mathcal{W}_E[\mathcal{A}] \cup \mathcal{W}_E[\mathcal{A}^c] \cup \tilde{J}^+[\mathcal{E}_A] \cup \tilde{J}^-[\mathcal{E}_A] \quad (5.1)$$

which is directly analogous to the boundary decomposition (2.1). However, for mixed states, where typically $\mathcal{E}_A \neq \mathcal{E}_{A^c}$, the decomposition (5.1) is *not* true;²⁶ instead the correct decomposition should replace $\mathcal{W}_E[\mathcal{A}^c]$ with the bulk domain of dependence of the complement of \mathcal{R}_A within the bulk Cauchy slice $\tilde{\Sigma}$, or more precisely $\tilde{D}[\tilde{\Sigma} \setminus \mathcal{R}_A \setminus \mathcal{E}_A]$.

Dual of ρ_A ? Within the class of CFTs and states with a geometrical holographic dual, it has often been asked,²⁷ for a given region \mathcal{A} , what is the bulk “dual” of the reduced density matrix ρ_A . One way to formulate what one means by this is as follows: suppose we fix ρ_A and vary over all compatible density matrices for the full state ρ . What is the maximal bulk spacetime region which coincides for all such ρ ’s? By

²⁶ Note however that if we purify a mixed state by additional boundaries, such as in the deformed eternal black hole example illustrated in Fig. 10, then the decomposition (5.1) does hold.

²⁷ In recent years this question has been invigorated by e.g. [53, 54].

“coinciding bulk regions” one means having the same geometry, i.e. the same bulk metric modulo diffeomorphisms. Another way to define the dual of $\rho_{\mathcal{A}}$ is to ask what is the maximal bulk region wherein we can uniquely determine the bulk metric (again modulo diffeomorphisms). In fact there are several (generally distinct) bulk regions that might be naturally associated with the density matrix; in nested order:

- The bulk region that $\rho_{\mathcal{A}}$ is *sensitive to*; in other words, regions wherein a deformation of the metric affects $\rho_{\mathcal{A}}$.²⁸
- The bulk region that $\rho_{\mathcal{A}}$ *determines*, i.e. where we can uniquely reconstruct all the components of the metric (up to diffeomorphisms).
- The bulk region that $\rho_{\mathcal{A}}$ *affects*, i.e. where by changing $\rho_{\mathcal{A}}$ one can change the bulk metric.

Here we focus on the second case, following [53, 54]. Based on lightsheet arguments, the authors of [53] proposed the causal wedge as the correct dual. On the other hand, [54], as well as [12, 22], argued that the requisite region should contain more than the causal wedge. In particular, [54] presented a number of criteria that such a region should satisfy, and explored several possibilities, most notably the region they denoted $\hat{w}(D_{\mathcal{A}})$ which corresponds to the bulk domain of dependence of the spacetime region spanned by all codimension-two extremal surfaces anchored within $D[\mathcal{A}]$. If every point of $\mathcal{R}_{\mathcal{A}}$ lies on at least one of these, then this region coincides with our entanglement wedge $\mathcal{W}_{\mathcal{E}}[\mathcal{A}]$. On the other hand, as [54] pointed out, there may be “holes” in such a set, i.e., regions of $\mathcal{R}_{\mathcal{A}}$ which do not lie along any least-area extremal surface anchored on a given region $\mathcal{A}' \subset \mathcal{A}$.²⁹

We propose that, since the most “natural” causal set associated with $\rho_{\mathcal{A}}$ from the bulk point of view is the entanglement wedge, this is indeed the most appropriate region to be identified with the “dual” of the reduced density matrix $\rho_{\mathcal{A}}$ (even in the presence of such entanglement “holes”). In this context, we should note that we can strip away the rest of the boundary spacetime, and consider the field theory just on $D[\mathcal{A}]$, which is a globally hyperbolic spacetime in its own right, in the state $\rho_{\mathcal{A}}$. Whether this state in general admits a holographic description is not known, but, if it does, then

²⁸ In fact there is a further subdivision here based on whether *any* geometrical deformation of the metric should change $\rho_{\mathcal{A}}$ or merely whether there should exist *some* deformation of the metric which changes $\rho_{\mathcal{A}}$. We thank Mark Van Raamsdonk for discussions on this issue.

²⁹ The example given in [54] involves a region through which traversing surfaces are not the smallest-area ones anchored on the given region, but a simpler physical example would be a point sufficiently close to an event horizon of an eternal spherical black hole, with $\mathcal{A} = \Sigma$ of one side as considered in §2.4.

a natural candidate would seem to be the entanglement wedge: this is, in its own right, a globally hyperbolic, asymptotically AdS spacetime, whose conformal boundary (according to theorem 13) is precisely $D[\mathcal{A}]$, and the area of whose edge $\mathcal{E}_{\mathcal{A}}$ gives the entropy of $\rho_{\mathcal{A}}$.

Here the word “natural” should be qualified, especially in light of the arguments in [22] that the causal wedge $\mathcal{W}_C[\mathcal{A}]$ is a natural bulk codimension-zero region associated with \mathcal{A} . The latter can be obtained more minimally: it suffices to know the causal structure of the bulk to define $\mathcal{W}_C[\mathcal{A}]$. On the other hand, the density matrix clearly encodes much more than the bulk causal structure, since at least it knows the entanglement entropy (as well as entanglement entropies of all subregions, apart from other observables). Since, in the bulk, the corresponding extremal surface is defined only once we know the bulk geometry, the entanglement wedge $\mathcal{W}_E[\mathcal{A}]$ it defines is a less minimal construct than the causal wedge $\mathcal{W}_C[\mathcal{A}]$. Nevertheless, once $\mathcal{E}_{\mathcal{A}}$ is identified, the rest of the bulk construction of the entanglement wedge is purely causal, and therefore defined fully robustly for any time-dependent asymptotically AdS spacetime.

The statement that the entanglement wedge is the natural dual of the reduced density matrix (which implies that the boundary observer in $D[\mathcal{A}]$ can learn about the bulk geometry in the entire $\mathcal{W}_E[\mathcal{A}]$) has a profound consequence. We have shown that the extremal surface $\mathcal{E}_{\mathcal{A}}$ has to lie in the causal shadow. This set can however be quite large, and so $\mathcal{E}_{\mathcal{A}}$ can lie very deep inside the bulk (as indicated by the shaded region in Fig. 10). In fact, a simple example supports the idea that the entanglement wedge represents the state in such a case (see Fig. 11). We start with a deconfined thermal state at $t = 0$ on a single \mathbf{S}^{d-1} , represented holographically by the exterior Schwarzschild-AdS solution. We add an outgoing null shell that reaches the boundary at $t < 0$ and an ingoing one that leaves it at $t > 0$. At $t = 0$ we still have the thermal state. The bulk solution is also unchanged between the past and future shells. However, these shells move the singularity and therefore have the effect of bringing the future and past event horizons closer to the boundary, leaving the previous bifurcation surface hidden behind both horizons. While this surface is no longer the bifurcation surface of a global Killing vector, it remains the extremal surface whose area gives the entropy of the state of the field theory on the right boundary. Presumably the holographic description of the state extends all the way down to this extremal surface, as it does in the absence of the shells, and thus consists of the entire entanglement wedge.

Another (related) example where the separation between entanglement wedge and causal wedge is particularly striking is the eternal (two-sided) black hole deformed by many shocks considered in [35, 55]. The Einstein-Rosen bridge is highly elongated and the extremal surface probably lies somewhere in the middle of it—so that the entanglement wedge for the entire right boundary is substantially larger than the causal




Fig. 11: Left: Exterior AdS-Schwarzschild solution, dual to a deconfined thermal state on S^{d-1} . The extremal surface for the entire boundary (red dot) coincides with the bifurcation surface and the causal information surface. Right: Vaidya solution with an ingoing null shell that reaches the boundary at $t < 0$ and an outgoing one that leaves it at $t > 0$ (brown); the geometry between the shells is unchanged, but the past and future event horizons (blue) have moved closer to the boundary, leaving the extremal surface (red dot) hidden behind them. The entanglement wedge in both cases is the entire spacetime (with a homology surface shown in green), while the causal wedge in the right figure is just the part outside of the event horizons. (The causal information surface is shown as the black dot.)

wedge, which in this case is simply the right exterior (domain of outer communication) of the black hole. So not only does the entanglement wedge penetrate arbitrarily close to the curvature singularity, it also contains a substantial part of the spacetime far beyond the black hole horizon!

Acknowledgments

We would like to thank Hong Liu, Juan Maldacena, Don Marolf, Steve Shenker, Mark Van Raamsdonk, and Aron Wall for discussions. MH, VH, and MR would like to thank University of Amsterdam for hospitality during the initial stages of this project. VH and MR would also like to acknowledge the hospitality of Brandeis University, MIT, Cambridge University, and the University of British Columbia, Vancouver during various

stages of this project. We would like to thank the Aspen Center for Physics (supported by the National Science Foundation under Grant 1066293) for their hospitality during both intermediate and concluding stages of this project.

MH is supported in part by the National Science Foundation under CAREER Grant No. PHY10-53842. VH and MR are supported in part by the Ambrose Monell foundation, by the FQXi grant “Measures of Holographic Information” (FQXi-RFP3-1334) and by the STFC Consolidated Grant ST/J000426/1. AL is supported in part by DOE grant DE-SC0009987. MR is supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013), ERC Consolidator Grant Agreement ERC-2013-CoG-615443: SPiN (Symmetry Principles in Nature).

References

- [1] S. Ryu and T. Takayanagi, “Aspects of holographic entanglement entropy,” *JHEP* **08** (2006) 045, [hep-th/0605073](#).
- [2] S. Ryu and T. Takayanagi, “Holographic derivation of entanglement entropy from AdS/CFT,” *Phys. Rev. Lett.* **96** (2006) 181602, [hep-th/0603001](#).
- [3] V. E. Hubeny, M. Rangamani, and T. Takayanagi, “A Covariant holographic entanglement entropy proposal,” *JHEP* **0707** (2007) 062, [arXiv:0705.0016 \[hep-th\]](#).
- [4] B. Swingle, “Entanglement Renormalization and Holography,” *Phys.Rev.* **D86** (2012) 065007, [arXiv:0905.1317 \[cond-mat.str-el\]](#).
- [5] M. Van Raamsdonk, “Comments on quantum gravity and entanglement,” [arXiv:0907.2939 \[hep-th\]](#).
- [6] M. Van Raamsdonk, “Building up spacetime with quantum entanglement,” *Gen.Rel.Grav.* **42** (2010) 2323–2329, [arXiv:1005.3035 \[hep-th\]](#).
- [7] J. Maldacena and L. Susskind, “Cool horizons for entangled black holes,” *Fortsch.Phys.* **61** (2013) 781–811, [arXiv:1306.0533 \[hep-th\]](#).
- [8] M. Headrick and T. Takayanagi, “A Holographic proof of the strong subadditivity of entanglement entropy,” *Phys.Rev.* **D76** (2007) 106013, [arXiv:0704.3719 \[hep-th\]](#).
- [9] P. Hayden, M. Headrick, and A. Maloney, “Holographic Mutual Information is Monogamous,” [arXiv:1107.2940 \[hep-th\]](#).
- [10] M. Headrick, “General properties of holographic entanglement entropy,” *JHEP* **1403** (2014) 085, [arXiv:1312.6717 \[hep-th\]](#).
- [11] A. Lewkowycz and J. Maldacena, “Generalized gravitational entropy,” [arXiv:1304.4926 \[hep-th\]](#).

- [12] A. C. Wall, “Maximin Surfaces, and the Strong Subadditivity of the Covariant Holographic Entanglement Entropy,” [arXiv:1211.3494 \[hep-th\]](#).
- [13] H. Casini, “Geometric entropy, area, and strong subadditivity,” *Class.Quant.Grav.* **21** (2004) 2351–2378, [arXiv:hep-th/0312238 \[hep-th\]](#).
- [14] S. Gao and R. M. Wald, “Theorems on gravitational time delay and related issues,” *Class.Quant.Grav.* **17** (2000) 4999–5008, [arXiv:gr-qc/0007021 \[gr-qc\]](#).
- [15] V. E. Hubeny, “Precursors see inside black holes,” *Int.J.Mod.Phys.* **D12** (2003) 1693–1698, [arXiv:hep-th/0208047 \[hep-th\]](#).
- [16] J. Abajo-Arrestia, J. Aparicio, and E. Lopez, “Holographic Evolution of Entanglement Entropy,” *JHEP* **1011** (2010) 149, [arXiv:1006.4090 \[hep-th\]](#).
- [17] V. E. Hubeny, “Extremal surfaces as bulk probes in AdS/CFT,” *JHEP* **1207** (2012) 093, [arXiv:1203.1044 \[hep-th\]](#).
- [18] T. Hartman and J. Maldacena, “Time Evolution of Entanglement Entropy from Black Hole Interiors,” *JHEP* **1305** (2013) 014, [arXiv:1303.1080 \[hep-th\]](#).
- [19] H. Liu and S. J. Suh, “Entanglement Tsunami: Universal Scaling in Holographic Thermalization,” *Phys.Rev.Lett.* **112** (2014) 011601, [arXiv:1305.7244 \[hep-th\]](#).
- [20] V. E. Hubeny and H. Maxfield, “Holographic probes of collapsing black holes,” [arXiv:1312.6887 \[hep-th\]](#).
- [21] V. E. Hubeny, M. Rangamani, and E. Tonni, “Thermalization of Causal Holographic Information,” [arXiv:1302.0853 \[hep-th\]](#).
- [22] V. E. Hubeny and M. Rangamani, “Causal Holographic Information,” *JHEP* **1206** (2012) 114, [arXiv:1204.1698 \[hep-th\]](#).
- [23] V. E. Hubeny, M. Rangamani, and E. Tonni, “Global properties of causal wedges in asymptotically AdS spacetimes,” *JHEP* **1310** (2013) 059, [arXiv:1306.4324 \[hep-th\]](#).
- [24] N. Engelhardt and A. C. Wall, “Quantum Extremal Surfaces: Holographic Entanglement Entropy beyond the Classical Regime,” [arXiv:1408.3203 \[hep-th\]](#).
- [25] R. Wald, *General Relativity*. University of Chicago Press, Chicago, 1984.
- [26] P. Buividovich and M. Polikarpov, “Entanglement entropy in gauge theories and the holographic principle for electric strings,” *Phys.Lett.* **B670** (2008) 141–145, [arXiv:0806.3376 \[hep-th\]](#).
- [27] W. Donnelly, “Decomposition of entanglement entropy in lattice gauge theory,” *Phys.Rev.* **D85** (2012) 085004, [arXiv:1109.0036 \[hep-th\]](#).
- [28] H. Casini, M. Huerta, and J. A. Rosabal, “Remarks on entanglement entropy for gauge fields,” *Phys.Rev.* **D89** (2014) 085012, [arXiv:1312.1183 \[hep-th\]](#).

- [29] W. Donnelly, “Entanglement entropy and nonabelian gauge symmetry,” [arXiv:1406.7304 \[hep-th\]](#).
- [30] D. Marolf, M. Rangamani, and T. Wiseman, “Holographic thermal field theory on curved spacetimes,” *Class.Quant.Grav.* **31** (2014) 063001, [arXiv:1312.0612 \[hep-th\]](#).
- [31] M. Headrick, “Entanglement Renyi entropies in holographic theories,” *Phys.Rev.* **D82** (2010) 126010, [arXiv:1006.0047 \[hep-th\]](#).
- [32] V. E. Hubeny, H. Maxfield, M. Rangamani, and E. Tonni, “Holographic entanglement plateaux,” *JHEP* **1308** (2013) 092, [arXiv:1306.4004](#).
- [33] V. E. Hubeny, “Covariant Residual Entropy,” [arXiv:1406.4611 \[hep-th\]](#).
- [34] V. Balasubramanian, P. Hayden, A. Maloney, D. Marolf, and S. F. Ross, “Multiboundary Wormholes and Holographic Entanglement,” [arXiv:1406.2663 \[hep-th\]](#).
- [35] S. H. Shenker and D. Stanford, “Black holes and the butterfly effect,” [arXiv:1306.0622 \[hep-th\]](#).
- [36] G. T. Horowitz and D. Marolf, “A New approach to string cosmology,” *JHEP* **9807** (1998) 014, [arXiv:hep-th/9805207 \[hep-th\]](#).
- [37] V. Balasubramanian, P. Kraus, A. E. Lawrence, and S. P. Trivedi, “Holographic probes of anti-de Sitter space-times,” *Phys.Rev.* **D59** (1999) 104021, [arXiv:hep-th/9808017 \[hep-th\]](#).
- [38] B. G. Carneiro da Cunha, “Inflation and holography in string theory,” *Phys.Rev.* **D65** (2002) 026001, [arXiv:hep-th/0105219 \[hep-th\]](#).
- [39] J. M. Maldacena, “Eternal black holes in anti-de Sitter,” *JHEP* **0304** (2003) 021, [arXiv:hep-th/0106112 \[hep-th\]](#).
- [40] G. Galloway, K. Schleich, D. Witt, and E. Woolgar, “Topological censorship and higher genus black holes,” *Phys.Rev.* **D60** (1999) 104039, [arXiv:gr-qc/9902061 \[gr-qc\]](#).
- [41] G. Galloway, K. Schleich, D. Witt, and E. Woolgar, “The AdS / CFT correspondence conjecture and topological censorship,” *Phys.Lett.* **B505** (2001) 255–262, [arXiv:hep-th/9912119 \[hep-th\]](#).
- [42] B. Freivogel, V. E. Hubeny, A. Maloney, R. C. Myers, M. Rangamani, *et al.*, “Inflation in AdS/CFT,” *JHEP* **0603** (2006) 007, [arXiv:hep-th/0510046 \[hep-th\]](#).
- [43] X. Dong, “Holographic Entanglement Entropy for General Higher Derivative Gravity,” *JHEP* **1401** (2014) 044, [arXiv:1310.5713 \[hep-th\]](#).
- [44] J. Camps, “Generalized entropy and higher derivative Gravity,” *JHEP* **1403** (2014) 070, [arXiv:1310.6659 \[hep-th\]](#).

- [45] J. Erdmenger, M. Flory, and C. Sleight, “Conditions on holographic entangling surfaces in higher curvature gravity,” *JHEP* **1406** (2014) 104, [arXiv:1401.5075 \[hep-th\]](#).
- [46] X. O. Camanho, J. D. Edelstein, J. Maldacena, and A. Zhiboedov, “Causality Constraints on Corrections to the Graviton Three-Point Coupling,” [arXiv:1407.5597 \[hep-th\]](#).
- [47] T. Faulkner, A. Lewkowycz, and J. Maldacena, “Quantum corrections to holographic entanglement entropy,” *JHEP* **1311** (2013) 074, [arXiv:1307.2892](#).
- [48] V. Balasubramanian and S. F. Ross, “Holographic particle detection,” *Phys.Rev.* **D61** (2000) 044007, [arXiv:hep-th/9906226 \[hep-th\]](#).
- [49] J. M. Maldacena, “Wilson loops in large N field theories,” *Phys.Rev.Lett.* **80** (1998) 4859–4862, [arXiv:hep-th/9803002 \[hep-th\]](#).
- [50] S.-J. Rey and J.-T. Yee, “Macroscopic strings as heavy quarks in large N gauge theory and anti-de Sitter supergravity,” *Eur.Phys.J.* **C22** (2001) 379–394, [arXiv:hep-th/9803001 \[hep-th\]](#).
- [51] J. Louko, D. Marolf, and S. F. Ross, “On geodesic propagators and black hole holography,” *Phys.Rev.* **D62** (2000) 044041, [arXiv:hep-th/0002111 \[hep-th\]](#).
- [52] M. Headrick, V. E. Hubeny, A. E. Lawrence, and M. Rangamani, “work in progress,”.
- [53] R. Bousso, S. Leichenauer, and V. Rosenhaus, “Light-sheets and AdS/CFT,” *Phys.Rev.* **D86** (2012) 046009, [arXiv:1203.6619 \[hep-th\]](#).
- [54] 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 \[hep-th\]](#).
- [55] S. H. Shenker and D. Stanford, “Multiple Shocks,” [arXiv:1312.3296 \[hep-th\]](#).