

Global properties of causal wedges in asymptotically AdS spacetimes

Veronika E. Hubeny,^a Mukund Rangamani^a and Erik Tonni^b

^aCentre for Particle Theory & Department of Mathematical Sciences,
Science Laboratories, South Road, Durham DH1 3LE, U.K.

^bSISSA and INFN,
via Bonomea 265, 34136, Trieste, Italy

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

ABSTRACT: We examine general features of causal wedges in asymptotically AdS spacetimes and show that in a wide variety of cases they have non-trivial topology. We also prove some general results regarding minimal area surfaces on the causal wedge boundary and thereby derive constraints on the causal holographic information. We go on to demonstrate that certain properties of the causal wedge impact significantly on features of extremal surfaces which are relevant for computation of holographic entanglement entropy.

KEYWORDS: Gauge-gravity correspondence, AdS-CFT Correspondence, Classical Theories of Gravity, Black Holes

ARXIV EPRINT: [1306.4324](https://arxiv.org/abs/1306.4324)

Contents

1	Introduction and summary	1
1.1	Causal constructions: a review	2
1.2	Preview of results	4
2	Topological structure of causal wedge	6
2.1	Global Schwarzschild-AdS _{d+1}	7
2.2	BTZ	15
2.3	Boosted black hole	16
2.4	Stars in AdS	19
3	General features of $\Xi_{\mathcal{A}}$ and $\chi_{\mathcal{A}}$	20
3.1	Summary of properties	20
3.2	Proofs of simple properties	21
4	Implications for bulk extremal surfaces	26
5	Discussion	28

1 Introduction and summary

One of the fundamental properties of a Lorentzian spacetime is its causal structure. In the context of the holographic AdS/CFT correspondence, the importance of causality was noted very early on [1] and it was realized that bulk AdS causality should be at the very least compatible with boundary causality. In fact in asymptotically AdS spacetimes, causal propagation through the bulk cannot be faster than propagation along the boundary as a consequence of the gravitational time-delay effect [2, 3]. Nevertheless, we do not understand the holographic dictionary well enough to pin-point a particular feature of the boundary field theory data that we can directly associate with the bulk causal structure. It is however clear that for any boundary field theory state which is described by a semi-classical bulk geometry, one can use certain observables of the field theory to probe aspects of the bulk causal structure, as has been noted in various contexts in the past, see e.g. [4–10].

Of interest to us is the set of bulk spacetime points which is naturally associated with a particular spatial region of the boundary field theory. This question has been tackled in a number of different ways in the recent past [11–14]. The various constructions described in these works can be divided into two classes: those that use regions bounded in the bulk by extremal surfaces inspired by the holographic entanglement entropy [15] and its covariant generalization [16], and those that put causal relations at the center stage. From these analyses and earlier works [17] the following picture emerges: a given spatial region

\mathcal{A} , whilst certainly being cognizant of the causal wedge associated with it in the bulk, is however able to access information of part of the geometry beyond the causal wedge.

Nevertheless, as noted in [11, 13], one might say that the most natural bulk spacetime region associated with a spatial region \mathcal{A} on the boundary is the causal wedge $\diamond_{\mathcal{A}}$ and associated quantities, since these are constructed solely using causal relations. Furthermore, the causally inspired constructions, as we shall see, serve to bound other observables, and in a certain sense the causal wedge is the bare minimum that the boundary field theory region should reproduce. Inspired by this observation, we undertake an examination of causal wedges in asymptotically AdS spacetimes (see [18, 19] for some previous observations). Despite being a simple exercise, it reveals rather interesting surprises; in fact we will demonstrate that the causal structure constrains other observables such as the entanglement entropy in a non-trivial fashion.

1.1 Causal constructions: a review

To set the stage for our discussion let us quickly recall some basic concepts relevant for the causal constructions. Consider a (d -dimensional) boundary spacetime foliated by a set of ($d - 1$ dimensional) Cauchy slices Σ_t , labeled by boundary time t . We will consider spacelike ($d - 1$ dimensional) regions $\mathcal{A}_t, \mathcal{B}_t, \dots \in \Sigma_t$ (generically we drop the subscript t for notational simplicity). The complement of a region \mathcal{A} will be denoted as \mathcal{A}^c .

In a nutshell, the *causal wedge* $\diamond_{\mathcal{A}}$, associated to a boundary region \mathcal{A} , is the set of bulk spacetime points which lie in both the future and the past of the boundary domain of dependence¹ $\diamond_{\mathcal{A}}$ for the region \mathcal{A} ,

$$\diamond_{\mathcal{A}} \equiv J^-[\diamond_{\mathcal{A}}] \cap J^+[\diamond_{\mathcal{A}}]. \tag{1.1}$$

In other words, the causal wedge consists of the set of spacetime events through which there exists a causal curve which starts and ends in $\diamond_{\mathcal{A}}$. The boundary of $\diamond_{\mathcal{A}}$ restricted to the bulk, denoted $\partial \diamond_{\mathcal{A}}$, consists of two² null surfaces $\partial_{\pm}(\diamond_{\mathcal{A}})$ which are generated by null geodesics; the outgoing null geodesics ending on the future boundary of $\diamond_{\mathcal{A}}$ generate $\partial_+(\diamond_{\mathcal{A}})$ and the ingoing ones from past boundary of $\diamond_{\mathcal{A}}$ generate $\partial_-(\diamond_{\mathcal{A}})$. These bulk co-dimension one null surfaces $\partial_+(\diamond_{\mathcal{A}})$ and $\partial_-(\diamond_{\mathcal{A}})$ intersect along a bulk co-dimension two spacelike surface $\Xi_{\mathcal{A}}$, which for reasons explained in [13] we dub the *causal information surface*. In other words,

$$\partial \diamond_{\mathcal{A}} = \partial_+(\diamond_{\mathcal{A}}) \cup \partial_-(\diamond_{\mathcal{A}}) \quad \text{and} \quad \Xi_{\mathcal{A}} = \partial_+(\diamond_{\mathcal{A}}) \cap \partial_-(\diamond_{\mathcal{A}}). \tag{1.2}$$

By construction, the surface $\Xi_{\mathcal{A}}$ is anchored on the entangling surface $\partial \mathcal{A}$ of the selected region \mathcal{A} , i.e. $\partial(\Xi_{\mathcal{A}}) = \partial \mathcal{A}$. For orientation, these constructs are illustrated in figure 1, for planar AdS (left) and global AdS (right). We emphasize that our construction is fully general and covariant, requiring only a spacetime geometry that allows us to define causal curves. For convenience we focus on a causal spacetime with a smooth metric.

¹ The domain of dependence $\diamond_{\mathcal{A}}$ corresponds to the boundary spacetime region whose physics is fully determined by initial conditions at \mathcal{A} . More formally, any fully extended timelike curve on the boundary which passes through $\diamond_{\mathcal{A}}$ must necessarily intersect \mathcal{A} .

² There may in fact be multiple boundaries; but we postpone a discussion of this subtlety till section 2.




Figure 1. A sketch of the causal wedge $\diamond_{\mathcal{A}}$ and associated quantities in planar AdS (left) and global AdS (right) in 3 dimensions taken from [20]: in each panel, the region \mathcal{A} is represented by the red curve on right, and the corresponding surface $\Xi_{\mathcal{A}}$ by blue curve on left; the causal wedge $\diamond_{\mathcal{A}}$ lies between the AdS boundary and the null surfaces $\partial_+(\diamond_{\mathcal{A}})$ (red surface) and $\partial_-(\diamond_{\mathcal{A}})$ (blue surface).

Given the bulk co-dimension two surface $\Xi_{\mathcal{A}}$, one can associate a scalar quantity with the region \mathcal{A} . In analogy with the definition of the holographic entanglement entropy [15, 16, 21], in [13] we defined the *causal holographic information* of \mathcal{A} , abbreviated $\chi_{\mathcal{A}}$, as quarter of the proper area of $\Xi_{\mathcal{A}}$ in Planck units,

$$\chi_{\mathcal{A}} \equiv \frac{\text{Area}(\Xi_{\mathcal{A}})}{4G_N}. \tag{1.3}$$

Although this number is infinite since $\Xi_{\mathcal{A}}$ reaches to the AdS boundary, as with entanglement entropy, there is information both in the divergence structure as well as in the regulated quantity.

The simplicity of these bulk constructs suggests that they should have correspondingly natural field theory dual. Moreover, the importance of causal structure in bulk gravity indicates that the dual field theory constructs should likewise correspond to fundamental quantities. In order to obtain hints of what these field theoretic quantities are, we set out to explore the bulk properties of the causal wedge and associated constructs. In [13] we have considered certain static geometries in order to study its properties in equilibrium, and suggested that $\chi_{\mathcal{A}}$ captures the basic amount of information about the bulk geometry contained in the reduced density matrix $\rho_{\mathcal{A}}$ for the region \mathcal{A} . While $\chi_{\mathcal{A}}$ provides a (rather weak upper) bound on the entanglement entropy $S_{\mathcal{A}}$, for certain special cases — which happen to be the ones where we can actually calculate the entanglement entropy from first principles — the causal information surface $\Xi_{\mathcal{A}}$ in fact coincides with the extremal surface $\mathfrak{E}_{\mathcal{A}}$, so this bound is saturated: $\chi_{\mathcal{A}} = S_{\mathcal{A}}$. To glean more intuition about the dynamical nature of our constructs, in [20] we focused specific time-dependent backgrounds corresponding to rapid thermalization. We observed that in general, if we consider a region \mathcal{A} at a certain boundary time $t_{\mathcal{A}}$, the construction of $\Xi_{\mathcal{A}}$ is quasi-teleological (i.e.

teleological only on light-crossing timescale) in time-dependent backgrounds, since the domain of dependence $\diamond_{\mathcal{A}}$ contains times later than $t_{\mathcal{A}}$.³ Finally, the recent analysis of [22] attempts to provide some interesting speculations about the field theoretic interpretation of $\chi_{\mathcal{A}}$.

In the present work, we continue the exploration of properties of causal wedge and related constructs, now in complete generality. This presents a somewhat complementary approach to that of [20]: instead of focusing on specific class of spacetimes where the explicit computation of $\chi_{\mathcal{A}}$ is tractable, we maximally relax the assumptions about the bulk spacetime, and consider global properties that these constructs, especially $\diamond_{\mathcal{A}}$ and $\Xi_{\mathcal{A}}$, must satisfy in general. Nevertheless, to exemplify our statements and familiarize the reader with the subtleties, we will present explicit constructions of $\Xi_{\mathcal{A}}$ in a particular spacetime in section 2.

1.2 Preview of results

We will start in section 2 by examining the topological structure of the causal wedge $\diamond_{\mathcal{A}}$ for a simply connected boundary region \mathcal{A} . A similar problem was considered in [19], with the conclusion that for simply connected regions \mathcal{A} , the causal wedge must likewise be simply connected. This follows from the statement of topological censorship [23], which asserts that every causal curve which begins and ends on the boundary of asymptotically AdS spacetime $\mathcal{I} = \partial\text{AdS}$ is contractible to the boundary. Viewing the domain of dependence $\diamond_{\mathcal{A}} \subset \mathcal{I}$ as a particular causal sub-domain, the result follows.

Here we will point out that the situation is actually more subtle. Despite the simplicity of \mathcal{A} , the associated causal wedge may be topologically arbitrarily complicated: for example, it can have (arbitrarily) many holes, i.e., non-trivial homology. In the examples we encounter, we will show that in asymptotically AdS_{d+1} bulk spacetime, it is possible for $\partial\diamond_{\mathcal{A}}$ to possess non-contractible co-dimension two spheres \mathbf{S}^{d-1} . Thus $\Xi_{\mathcal{A}}$ may be composed of multiple disconnected components. Moreover, the change in topology can be engineered by varying parameters associated with \mathcal{A} relative to those characterizing the bulk geometry, implying interesting ‘phase transitions’ for χ .

This can be observed already in perhaps the simplest non-trivial example distinct from pure AdS, namely a neutral static black hole in global AdS_{d+1} for $d > 2$, provided \mathcal{A} is large enough, as explained in section 2.1 (cf. figure 3 for an illustrative plot). Despite the simplicity of the setup, this result may come as a surprise to many readers, and in fact (to our knowledge) has not been pointed out previously in the literature. There may be several reasons for this. Initial studies of the causal wedge (such as explicitly carried out in [20]) often focus on 3-dimensional bulk geometry, where this effect is absent. For example, the BTZ black hole causal wedge is simply connected for any region size, as we explicitly illustrate in section 2.2. Moreover, most higher-dimensional studies typically focus on planar black holes, so as to consider duals of states of CFT on Minkowski background.

³ Curiously, while χ behaves correspondingly quasi-teleologically in general time-dependent backgrounds, in the case of collapsing thin null shell examined in [20], the temporal evolution of χ remains entirely causal. In other words, while the causal information surface Ξ gets deformed by the shell quasi-teleologically, its area χ remains unaffected till after the region \mathcal{A} itself encounters the shell.

However, the effect we describe is absent for a planar black hole, since there is no region ‘on the other side’ of the black hole which could be causally accessible from the boundary. We illustrate explicitly that when the CFT lives on the Einstein static universe, $\Xi_{\mathcal{A}}$ is disconnected whenever the region \mathcal{A} covers a sufficient portion of the full system (ranging from more than half of the full system in case of tiny black holes to almost the entire system in case of huge black holes).

Considering large regions \mathcal{A} (i.e. comparable to the size of the full system on compact space) may perhaps seem too artificial to worry about. To dispel such objections, we proceed in section 2.3 to argue that in fact the causal wedge for arbitrarily small region can also have holes, for a suitably chosen bulk geometry. A simple example is provided by the global completion of the conformal soliton [24], which is in fact just a boosted version of a static black hole [1]. One can also consider genuinely dynamical situations involving multiple black holes. Though the explicit metric is not known analytically, such configurations provide simple existence arguments for causal wedges with multiple holes (and correspondingly multiple components of a disconnected $\Xi_{\mathcal{A}}$).

Based on the above examples, one might easily wonder if the non-trivial structure of a causal wedge is somehow inherited from the bulk geometry being causally non-trivial. It is in fact easy to argue that this is *not* the case, as we discuss in section 2.4. A causally trivial spacetime, such as a compact star, can likewise admit causal wedges with holes. This further solidifies the robustness of this feature.

Having seen in section 2 that even in simple bulk spacetimes causal wedges may have surprising properties, the reader might be led to suspect that this will render the CFT dual of the causal wedge and derived quantities far too complicated. However, we take the viewpoint that since the causal wedge $\blacklozenge_{\mathcal{A}}$ is the simplest and most natural bulk construct associated with a boundary region \mathcal{A} , it ought to have a natural CFT dual nevertheless. In section 3 we collect the global properties that any causal wedge must satisfy. These include natural inclusion properties, as well as simple additivity properties of the causal holographic information χ , which are previewed in section 3.1 and justified in section 3.2.

Having established these global properties for the constructs ($\blacklozenge_{\mathcal{A}}$, $\Xi_{\mathcal{A}}$, and $\chi_{\mathcal{A}}$) derived from our causal wedge, in section 4 we turn to a brief discussion of implications specifically for the extremal surface $\mathfrak{E}_{\mathcal{A}}$ and the entanglement entropy $S_{\mathcal{A}}$ associated to a given boundary region \mathcal{A} . Most intriguingly, the property established in section 3 that any extremal surface $\mathfrak{E}_{\mathcal{A}}$ must lie outside⁴ of the causal wedge \mathcal{A} , together with the observation of section 2 that the causal wedge has a hole (i.e. Ξ consists of two disconnected components) in the Schwarzschild-AdS background for sufficiently large region \mathcal{A} , implies that for such cases, there does not exist a connected extremal surface anchored on $\partial\mathcal{A}$ homologous to \mathcal{A} . In particular, the connected extremal surface $\mathfrak{E}_{\mathcal{A}^c}$ corresponding to the complement \mathcal{A}^c of the region \mathcal{A} does not satisfy the homology requirement since it is separated from \mathcal{A} by the black hole, while a surface going around the black hole so as to be homologous to \mathcal{A} would necessarily have to enter inside the causal wedge. This means that the extremal surface whose area gives the entanglement entropy $\mathfrak{E}_{\mathcal{A}}$ must likewise consist of two dis-

⁴ By outside we mean here that no part of the extremal surface $\mathfrak{E}_{\mathcal{A}}$ can lie within the causal wedge $\blacklozenge_{\mathcal{A}}$.

connected components: one given by $\mathfrak{E}_{\mathcal{A}^c}$ and the other wrapping the black hole horizon. That in turn implies that for such cases (i.e. sufficiently large \mathcal{A}), the difference in entanglement entropies $S_{\mathcal{A}} - S_{\mathcal{A}^c}$ captures precisely the thermal entropy S_{BH} , and therefore saturates the Araki-Lieb inequality [25].⁵ This is perhaps the most interesting result of our explorations, justifying our intuition that the properties of the causal wedges serve to non-trivially constrain physically understood observables such as entanglement entropy.

Given the potentially profound implications of the existence of causal wedges with holes, as well as the relative ease with which one can construct examples of such an occurrence, one might start to worry that perhaps for any bulk geometry different from pure AdS (and satisfying the null energy condition), we could construct a sufficiently large boundary region \mathcal{A} whose causal wedge $\diamond_{\mathcal{A}}$ has a hole. We address this possibility in section 5 and suggest that this does not happen unless the deformation is sufficiently strong; in spherically symmetric situations, we conjecture that the presence of holes in the causal wedge is associated with the presence of null circular orbits in the spacetime.

2 Topological structure of causal wedge

Let us first examine the topological structure of the causal wedge $\diamond_{\mathcal{A}}$ for a given bulk spacetime and a simply connected boundary region \mathcal{A} . Recall that the causal wedge can be thought of as consisting of causal curves which begin and end in $\diamond_{\mathcal{A}}$. Although these curves are all continuously deformable into each other, we show that the causal wedge can be topologically more complicated. For example, $\Xi_{\mathcal{A}}$ need not be homologous to \mathcal{A} , or even when it is, $\partial\diamond_{\mathcal{A}}$ need not be homotopic to $\diamond_{\mathcal{A}}$. We focus on $\Xi_{\mathcal{A}}$, specifically whether or not it is connected and the nature of transitions between the number of its components.

To motivate the possibility of $\Xi_{\mathcal{A}}$ having multiple components, let us observe that a causal wedge cannot penetrate a black hole (i.e., $\Xi_{\mathcal{A}}$ cannot reach beyond an event horizon). Heuristically this follows directly from the definition of a black hole: no causal curve from inside can reach the AdS boundary, much less $\diamond_{\mathcal{A}} \subset \mathcal{I}$ (see also section 3). Suppose however that the black hole is very small (compared to AdS scale). Far away from the black hole, its gravitational effects are negligible; so there is a spatial region surrounding the black hole from which causal curves which would have reached well within $\diamond_{\mathcal{A}}$ in pure AdS still reach $\diamond_{\mathcal{A}}$ in the actual spacetime. In other words, although the black hole itself cannot lie inside the causal wedge, a spatial region fully surrounding it is contained in $\diamond_{\mathcal{A}}$. This reasoning suggests that for small black holes and sufficiently large regions \mathcal{A} , the causal information surface $\Xi_{\mathcal{A}}$ has two components: one similar to that in pure AdS which is anchored on $\partial\mathcal{A}$, and one which shields the black hole. We will now check this expectation explicitly for Schwarzschild-AdS _{$d+1$} for $d \geq 3$ (focusing on $d = 4$ which is algebraically simplest) and then comment on other geometries.

⁵ This observation is further discussed in [26] where it is explicitly demonstrated that connected minimal surfaces fail to capture the entanglement entropy for large enough regions \mathcal{A} . The resulting saturation of the Araki-Lieb inequality is referred to there as the entanglement plateau phenomenon.

2.1 Global Schwarzschild-AdS_{d+1}

To illustrate the point that even for simply-connected regions \mathcal{A} the causal information surface $\Xi_{\mathcal{A}}$ may be composed of multiple disjoint components, let us consider the Schwarzschild-AdS_{d+1} black hole. We will w.l.o.g. set the AdS scale to unity, and characterize the black hole by its horizon size in AdS units, $r_h \in (0, \infty)$. This gives a 1-parameter family of static, spherically symmetric and physically well-behaved spacetimes with metric

$$ds^2 = -f(r) dt^2 + \frac{dr^2}{f(r)} + r^2 (d\varphi^2 + \sin^2 \varphi d\Omega_{d-2}^2), \quad f(r) = r^2 + 1 - \frac{r_h^{d-2} (r_h^2 + 1)}{r^{d-2}}. \quad (2.1)$$

Large black holes ($r_h > 1$) are dual to the thermal density matrix of the field theory on the Einstein Static Universe $\text{ESU}_d = \mathbf{S}^{d-1} \times \mathbb{R}$ (e.g. as in the case of $\mathcal{N} = 4$ SYM in $d = 4$), while small black holes are still physically relevant in the microcanonical ensemble.⁶ For convenience we will restrict attention to boundary regions \mathcal{A} which preserve $\text{SO}(d-1)$ spherical symmetry. For purposes of finding the causal wedge $\blacklozenge_{\mathcal{A}}$, we can then reduce this problem to effectively 3-dimensional one⁷ by reducing the \mathbf{S}^{d-1} to one non-trivial angle $\varphi \in [0, \pi]$. The region \mathcal{A} is then characterized by its radius $\varphi_{\mathcal{A}}$. As $\varphi_{\mathcal{A}} \rightarrow \pi$, the region covers most of the boundary space.⁸ On the other hand, the planar black hole case is recovered in the limit $\varphi_{\mathcal{A}} \rightarrow 0$ and $r_h \rightarrow \infty$.

It is convenient to use coordinates (t, ρ, φ) where $\rho \in [0, \pi/2)$ is related to the standard radial coordinate r by $r = \tan \rho$. The relevant 3-dimensional piece of the bulk metric is then

$$ds^2 = \frac{1}{\cos^2 \rho} \left(-g(\rho) dt^2 + \frac{d\rho^2}{g(\rho)} + \sin^2 \rho d\varphi^2 \right), \quad (2.2)$$

where

$$g(\rho) = 1 - \mu \frac{\cos^d \rho}{\sin^{d-2} \rho}, \quad \mu \equiv r_h^{d-2} (r_h^2 + 1) = \frac{\sin^{d-2} \rho_h}{\cos^d \rho_h}. \quad (2.3)$$

Null geodesics in this subspace are characterized by the reduced angular momentum $\ell \in (-1, 1)$, a discrete parameter $\eta = \pm 1$ labeling outgoing ($\eta = 1$) versus ingoing ($\eta = -1$) geodesics, as well as the initial position. We can write the differential equations in terms of the affine parameter λ where $\dot{} \equiv \frac{d}{d\lambda}$ as follows:

$$\dot{t} = \frac{\cos^2 \rho}{g(\rho)}, \quad \dot{\varphi} = \ell \frac{\cos^2 \rho}{\sin^2 \rho}, \quad \dot{\rho} = \eta \cos^2 \rho \sqrt{1 - \ell^2 \frac{g(\rho)}{\sin^2 \rho}}. \quad (2.4)$$

⁶ The translation between field theory and geometry is the following: the black holes with horizon size r_h have a Hawking temperature $T_{BH} = \frac{d r_h^2 + (d-2) \ell_{\text{AdS}}^2}{4\pi r_h \ell_{\text{AdS}}^2}$. These solutions have a minimum value of T_{BH} attained at $r_h = \sqrt{\frac{d-2}{d}} \ell_{\text{AdS}}$. They however minimize the free energy only for $T_{BH} \ell_{\text{AdS}} \geq \frac{d-1}{2\pi}$ or equivalently $r_h \geq \ell_{\text{AdS}}$.

⁷ Note that any curve which is causal in the full $(d+1)$ -dimensional space is necessarily causal in the reduced $(2+1)$ -dimensional subspace, and conversely any causal curve in the 3-dimensional spacetime trivially lifts to a causal curve in the full $(d+1)$ -dimensional spacetime.

⁸ To keep $\blacklozenge_{\mathcal{A}}$ finite, we will however consider $\varphi_{\mathcal{A}} < \pi$ in this section. (As explained in section 3, when $\varphi_{\mathcal{A}} = \pi$ the boundary region $\mathcal{A} = \mathbf{S}^{d-1}$ is a complete Cauchy slice of the Einstein Static Universe. Then $\blacklozenge_{\mathcal{A}}$ is simply the region exterior of the black hole and $\Xi_{\mathcal{A}}$ is the bifurcation surface of the horizon.)

We could in principle obtain analytic expressions for $(t(\lambda), \rho(\lambda), \varphi(\lambda))$, which are given in terms of elliptic functions. This however does not add much insight and it is easier to see the structure graphically, so we simply integrate the geodesic equations numerically. In actual implementation it is convenient to solve (2.4) for $t(\rho)$ and $\phi(\rho)$ directly, though we have to keep track of η changing sign at a turning point where $\dot{\rho} = 0$.

Only geodesics with sufficiently large angular momentum have a turning point; these are ones for which the equation $\dot{\rho} = 0$ in (2.4) has a real solution $\rho_0 \in (0, \frac{\pi}{2})$. It is easy to check that this only occurs for $\ell^2 \in (\ell_0^2, 1)$, where

$$\ell_0 = \left[1 + (d-2) d^{-\frac{d}{d-2}} \left(\frac{2}{\mu} \right)^{\frac{2}{d-2}} \right]^{-1/2} \xrightarrow{d=4} \sqrt{\frac{4\mu}{1+4\mu}} = \frac{2 r_h \sqrt{r_h^2 + 1}}{2 r_h^2 + 1}. \quad (2.5)$$

The corresponding value of ρ_0 at this minimal ℓ_0 corresponds to the circular null orbit radius and is given by

$$\rho_0(\ell_0) = \tan^{-1} \left(\frac{d\mu}{2} \right)^{\frac{1}{d-2}} \xrightarrow{d=4} \tan^{-1} \sqrt{2\mu}. \quad (2.6)$$

For general $\ell > \ell_0$, the radial position of the turning point is given by the largest root of the polynomial $\dot{\rho} = 0$ which is of order d (or $\frac{d}{2}$ for even d); e.g.

$$\rho_0 = \tan^{-1} \sqrt{\frac{\ell^2}{2(1-\ell^2)} \left[1 + \sqrt{1 - 4\mu \left(\frac{1-\ell^2}{\ell^2} \right)} \right]} \quad \text{for } d = 4. \quad (2.7)$$

(On the other hand, geodesics with $\ell < \ell_0$ have no turning point: instead they terminate at the curvature singularity at $r = 0$.)

Now that we have ρ_0 , we can find the corresponding φ and t coordinates for the turning point. We typically specify the particular ℓ -geodesic by giving the functions $t_\ell(\rho)$ and $\varphi_\ell(\rho)$ and when necessary explicitly indicating whether we are before or after the turning point. In particular, we have the following expressions for the trajectory of the future-directed null geodesics from a boundary point q^\vee defined below (hence initially $\eta = -1$) parameterized by ℓ :

$$t_\ell(\rho) = t_i + \int_{\rho_i}^{\rho_f} \frac{h_\ell(\tilde{\rho})}{g(\tilde{\rho})} d\tilde{\rho}, \quad \varphi_\ell(\rho) = \varphi_i + \ell \int_{\rho_i}^{\rho_f} \frac{h_\ell(\tilde{\rho})}{\sin^2 \tilde{\rho}} d\tilde{\rho}, \quad h_\ell(\rho) \equiv \frac{1}{\sqrt{1 - \ell^2 \frac{g(\rho)}{\sin^2 \rho}}},$$

ingoing segment ($\eta = -1$): $t_i = t_\vee, \quad \varphi_i = 0, \quad \rho_i = \rho, \quad \rho_f = \frac{\pi}{2},$

outgoing segment ($\eta = 1$): $t_i = t_{\text{ingoing}}(\rho_0) \equiv t_0, \quad \varphi_i = \varphi_{\text{ingoing}}(\rho_0) \equiv \varphi_0, \quad \rho_i = \rho_0, \quad \rho_f = \rho.$

(2.8)

Similar expressions can be written down for the geodesics that end up on q^\wedge .

Let us now turn to the strategy for finding $\Xi_{\mathcal{A}}$. Since the geometry is static, we can w.l.o.g. place \mathcal{A} at time $t = 0$. The domain of dependence $\diamond_{\mathcal{A}}$ for round regions of radius $\varphi_{\mathcal{A}}$ is then determined by two boundary points $\diamond_{\mathcal{A}} = J^+[q^\vee] \cap J^-[q^\wedge]$ with (t, ρ, φ) coordinates $q^\vee = (-\varphi_{\mathcal{A}}, \frac{\pi}{2}, 0)$ and $q^\wedge = (\varphi_{\mathcal{A}}, \frac{\pi}{2}, 0)$.⁹ Recall that for general spacetimes the boundary of

⁹ We adapt the notation introduced in [13]; for simple regions the causal wedge is generated by null geodesics emanating from two points q^\wedge in the future and q^\vee in the past.

the full causal wedge $\diamond_{\mathcal{A}}$ is generated by bulk null geodesics which start at q^\vee (for $\partial_-(\diamond_{\mathcal{A}})$) or end on q^\wedge (for $\partial_+(\diamond_{\mathcal{A}})$). Their intersection gives $\Xi_{\mathcal{A}}$. Since the Schwarzschild-AdS spacetime is static, the two sets of geodesics are merely time-reversed versions of each other, and their intersection necessarily lies at $t = 0$. Furthermore, by spherical symmetry, each congruence respects φ -reversal symmetry. This means that it suffices to find just one congruence to determine the rest. For convenience of discussion let us label the four sets of geodesics by letters P (F) for past (future) congruence and R (L) for right (left) part of each congruence, i.e. ingoing from the boundary towards positive (negative) φ . So PR and PL geodesics generate $\partial_-(\diamond_{\mathcal{A}})$ and FR and FL geodesics generate $\partial_+(\diamond_{\mathcal{A}})$, with $\ell > 0$ along the PR and FL congruence. Hence we only need to find the solution to (2.4) given by $t_\ell(\rho)$ and $\varphi_\ell(\rho)$ for the PR congruence, say; the symmetries

$$PR_\ell(t, \rho, \varphi) = PL_\ell(t, \rho, -\varphi) = FR_\ell(-t, \rho, \varphi) = FL_\ell(-t, \rho, -\varphi) \quad (2.9)$$

then immediately give the PL , FR , and FL congruences.

Let us now consider the intersections of these congruences.¹⁰ For each ℓ , PR_ℓ geodesic intersects FR_ℓ geodesic at $t = 0$ and similarly for the PL_ℓ and FL_ℓ geodesics. Let us denote the curve generated by these intersections (i.e. parameterized by ℓ) on the $t = 0$ surface by $\mathcal{X}_{t=0}$, and let us denote its coordinates by $(\rho_{t=0}(\ell), \varphi_{t=0}(\ell))$ and $(\rho_{t=0}(\ell), -\varphi_{t=0}(\ell))$ for the R and L halves of the congruences respectively. For each ℓ , we can find these by first solving $t_\ell(\rho) = 0$ for ρ and then substituting this into $\varphi_\ell(\rho)$ to determine φ . Consider now the function $\varphi_{t=0}(\ell)$. For the radial geodesics, we necessarily have $\varphi_{t=0}(0) = 0$, whereas for the boundary geodesics, it is easy to see that $\varphi_{t=0}(1) = \varphi_{\mathcal{A}} < \pi$. If $\varphi_{t=0}(\ell)$ increases monotonically with ℓ , or more generally if $\varphi_{t=0}(\ell) < \pi$ for all $\ell \in (0, 1)$, then $\mathcal{X}_{t=0}$ gives the full curve $\Xi_{\mathcal{A}}$.

However, there exists another set of intersections between the congruences which becomes relevant if $\varphi_{t=0}(\ell) > \pi$. In particular, the FR_ℓ geodesic intersects the FL_ℓ geodesic at $\varphi = \pi$, and similarly for PR and PL geodesics. When $\varphi_{t=0}(\ell) > \pi$, the PR geodesic from q^\vee (along which φ starts from 0 and increases monotonically with the affine parameter) would have intersected the PL geodesic (at $\varphi = \pi$) before intersecting the FR geodesic (at $t = 0$ and $\varphi = \varphi_{t=0}(\ell) > \pi$). The moment two null geodesics from q^\vee intersect,¹¹ they henceforth become timelike-separated from q^\vee , and therefore enter inside $\diamond_{\mathcal{A}}$, no longer remaining on $\partial_-(\diamond_{\mathcal{A}})$. The subsequent intersection on $\mathcal{X}_{t=0}$ is therefore not on $\partial_-(\diamond_{\mathcal{A}})$ and correspondingly is not relevant for $\Xi_{\mathcal{A}}$. Said differently, $\Xi_{\mathcal{A}}$ closes off at $\varphi = \pi$.

To summarize, the condition for $\Xi_{\mathcal{A}}$ to have two disconnected components is

$$\max_{\ell \in (0,1)} \varphi_{t=0}(\ell) > \pi . \quad (2.10)$$

If (2.10) holds, then there are two¹² solutions of $\varphi_{t=0}(\ell) = \pi$; let us label them by $0 < \ell_1 < \ell_2 < 1$. In such a case, $\Xi_{\mathcal{A}}$ has one component given by $\mathcal{X}_{t=0}$ for $\ell \in (0, \ell_1)$ and

¹⁰ For the Schwarzschild-AdS_{d+1} spacetime, it can be checked that geodesics *within* each congruence do not intersect each other.

¹¹ In slight abuse of language but following previous terminology, we will refer to these intersections as caustics and denote them by \mathcal{C}^\pm for the $\partial_\pm(\diamond_{\mathcal{A}})$ congruences.

¹² A-priori, there could have been an even number larger than 2, but explicit checks indicate that this




Figure 2. A plot of the intersection points of the future and past congruence, $\mathcal{X}_{t=0}$, plotted on the Poincaré disk. In each panel, the outer circle represents the AdS boundary (with the region \mathcal{A} highlighted in red; $\varphi_{\mathcal{A}} = 2.5$ in both panels). The black hole size is $r_h = 0.5$ (left) and $r_h = 0.2$ (right), denoted by red dashed curve (but obscured in the latter case). $\mathcal{X}_{t=0}$ is composed of the individual intersection points, color-coded by ℓ (from red at $\ell = 0$ to purple at $\ell = 1$). For large enough black hole (left), $\varphi_{t=0}(\ell) < \pi$ for all ℓ , and therefore $\mathcal{X}_{t=0} = \Xi_{\mathcal{A}}$. For small black hole (right) $\mathcal{X}_{t=0}$ self-intersects and therefore $\Xi_{\mathcal{A}}$ has two components as indicated.

another given by $\mathcal{X}_{t=0}$ for $\ell \in (\ell_2, 1)$. The latter is connected to the AdS boundary and is anchored at $\partial\mathcal{A}$ for $\ell = 1$. The former is disconnected from the boundary and wraps the black hole. This situation is illustrated in figure 2, where we plot $\mathcal{X}_{t=0}$ on the Poincaré disk for connected (left) and disconnected (right) case.

To understand better what happens in the disconnected case, it is instructive to consider the full causal wedge. This is illustrated in figure 3 where we plot the causal wedge for the same set of parameters as in the right panel of figure 2, but now on a 3-d spacetime diagram in ingoing Eddington coordinates.¹³ There are several features of note: as expected, the causal wedge clearly has a hole, causing $\Xi_{\mathcal{A}}$ to have two disconnected components, one anchored on $\partial\mathcal{A}$ and one wrapping the black hole. This was already necessitated by the observation that the causal wedge cannot penetrate the black hole, while approximating the pure AdS causal wedge far away from the black hole. However, unlike the pure AdS

doesn't happen; essentially there isn't enough structure in the geodesic equations for Schwarzschild-AdS to allow multiple extrema. Said differently, there are two competing effects which influence how much a geodesic 'orbits' in a given time span: the light bending gets stronger nearer to the black hole, but so does the time-delay. To maximize the former while minimizing the latter, we need to tune ℓ to attain the optimal penetration depth (approximately given by the null circular orbit radius); our assertion follows since $\varphi_{t=0}(\ell)$ increases for smaller ℓ and decreases for larger ℓ .

¹³ Following previous convention [20], we plot ρ radially and choose the vertical coordinate such that ingoing radial null geodesics lie at 45° . This fixes the vertical coordinate to be given by $v - \rho + \frac{\pi}{2}$, where

$$v = t + \frac{1}{2r_h^2 + 1} \left(\sqrt{r_h^2 + 1} \left[\tan^{-1} \frac{\tan \rho}{\sqrt{r_h^2 + 1}} - \frac{\pi}{2} \right] - r_h \tanh^{-1} \frac{r_h}{\tan \rho} \right).$$

For this reason, the plot is asymmetric under vertical flip.

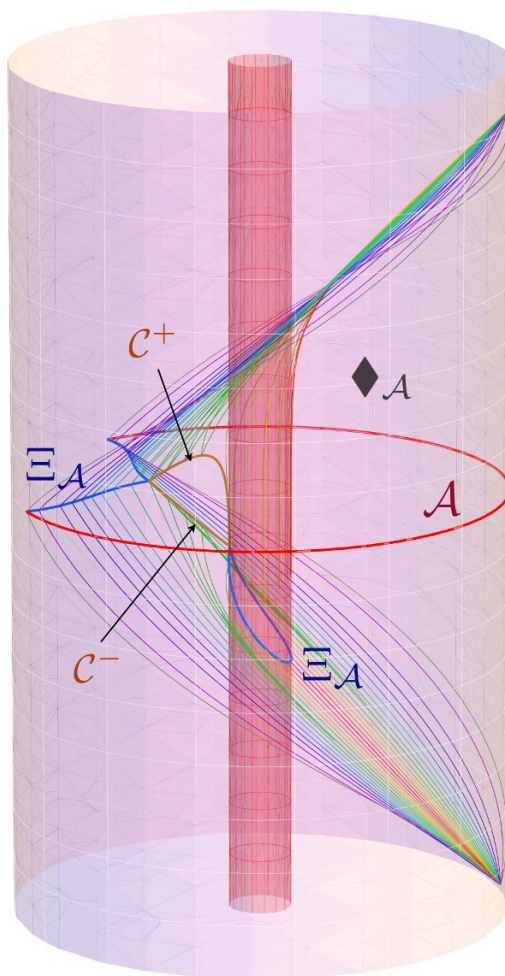


Figure 3. Causal wedge for the case $r_h = 0.2$ and $\varphi_{\mathcal{A}} = 2.5$, as in right panel of figure 2. Same color-coding (by ℓ) is applied to the null geodesic generators of $\partial_{\pm}\diamond_{\mathcal{A}}$. In addition to the AdS boundary and horizon, the plot exhibits the region \mathcal{A} (indicated by the thick red curve), the two components of $\Xi_{\mathcal{A}}$ (indicated by the thick blue curves), and the curves of caustics \mathcal{C}^{\pm} (indicated by thick brown curves) which connect up the two components of $\Xi_{\mathcal{A}}$. The causal wedge $\diamond_{\mathcal{A}}$ bounded by the null generators clearly exhibits a hole.

case, the boundary of the causal wedge has caustics where the L and R geodesics from the same congruence intersect each other (before intersecting those from the other congruence). The two caustic curves (\mathcal{C}^+ on $\partial_+(\diamond_{\mathcal{A}})$ and \mathcal{C}^- on $\partial_-(\diamond_{\mathcal{A}})$) lie at $\varphi = \pi$ and connect the two components of $\Xi_{\mathcal{A}}$, where the latter cusps.

As an aside, we remark that the presence of caustics in the causal wedge implies that generically the causal information surface $\Xi_{\mathcal{A}}$ need not be smooth; for less symmetric spacetimes this can happen even when $\Xi_{\mathcal{A}}$ is connected. Although that might seem like a bizarre feature, it is worth remembering that it is actually no worse than the analogous




Figure 4. The projection of the causal information surfaces $\Xi_{\mathcal{A}}$ for various $\varphi_{\mathcal{A}}$ onto the Poincaré disk of Schwarzschild-AdS₅ with fixed black hole size $r_h = 0.2$, color-coded by $\varphi_{\mathcal{A}}$ which varies from 0 (red) to π (purple) in increments of 0.1. (For example, the blue curve with $\varphi_{\mathcal{A}} = 2.5$ corresponds to the projection of $\Xi_{\mathcal{A}}$ in figure 3.) We can clearly see that $\Xi_{\mathcal{A}}$ pinches off; for larger $\varphi_{\mathcal{A}}$, the disconnected component of $\Xi_{\mathcal{A}}$ is located very near the horizon.

property of an event horizon, which likewise is not smooth generically.¹⁴ In particular, although a horizon generator has to remain on the horizon, new generators can enter the horizon at caustics.

Let us now return to considering the effect of varying the parameters, namely the transition between connected and disconnected $\Xi_{\mathcal{A}}$ exemplified in figure 2. As we decrease the black hole size r_h (for a fixed $\varphi_{\mathcal{A}}$), the curve $\varphi_{t=0}(\ell)$ reaches higher and higher, eventually exceeding π (and growing without bound as $r_h \rightarrow 0$). This is because the null geodesics whose angular momentum is close to ℓ_0 orbit around the black hole many ($\sim r_h^{-1}$) times before reaching $t = 0$. Conversely, for fixed r_h , the causal information surface $\Xi_{\mathcal{A}}$ is connected if $\varphi_{\mathcal{A}}$ is sufficiently small (and is guaranteed to be so if $\varphi_{\mathcal{A}} < \pi/2$ for any r_h) and disconnected if $\varphi_{\mathcal{A}}$ is large enough (and is guaranteed to be so when $\varphi_{\mathcal{A}} \rightarrow \pi$ for any r_h). This is illustrated in figure 4, which plots the causal information surfaces $\Xi_{\mathcal{A}}$ for the full range of $\varphi_{\mathcal{A}}$ at a fixed black hole size ($r_h = 0.2$, as in figure 3), projected onto the Poincaré disk. This presents a somewhat complementary information to that in figure 2: whereas the latter varied r_h at fixed $\varphi_{\mathcal{A}}$, figure 4 varies $\varphi_{\mathcal{A}}$ at fixed r_h .

Since $\varphi_{\mathcal{A}}$ is connected for small $\varphi_{\mathcal{A}}$ and disconnected for sufficiently large $\varphi_{\mathcal{A}}$, we're

¹⁴ Event horizons are locally Lipschitz, but not necessarily more regular than that [27]. There even exist (rather exotic) examples [28] where the event horizon is ‘nowhere’ differentiable, in the sense of non-existence of any open neighbourhood where the horizon is differentiable, though they are differentiable ‘almost everywhere’ in a measure-theoretic sense. More generally, [29] showed that differentiability fails precisely where new generators join the horizons.

guaranteed by continuity that for any r_h , there is a critical region size $\varphi_{\mathcal{A}}^*$ for which Ξ just pinches off. Let us denote the angular momentum along the corresponding null geodesic (i.e. the one passing through this pinch-off point) by ℓ_* . The four geodesics, PR_{ℓ_*} , PL_{ℓ_*} , FR_{ℓ_*} , and FL_{ℓ_*} , all intersect at a single point $t = 0, \varphi = \pi$, and $\rho = \rho_*$. While this is true whenever $\varphi_{t=0}(\ell) = \pi$, here we have an extra condition that $\Xi_{\mathcal{A}}$ self-intersects at a tangent, i.e. that $\max_{\ell \in (0,1)} \varphi_{t=0}(\ell) = \pi$. This condition is satisfied only for a specific relation between r_h and $\varphi_{\mathcal{A}}$. Hence to determine ℓ_* and ρ_* , and correspondingly the critical curve in $(r_h, \varphi_{\mathcal{A}})$ plane, we need to be able to find $\max_{\ell \in (0,1)} \varphi_{t=0}(\ell)$ efficiently.

While doing this numerically is rather time-consuming, we can simplify matters using the following observation. Geodesics with angular momentum ℓ_0 , when close to $\rho \approx \rho_0$, move very slowly in the ρ direction. Hence to compensate, i.e. to remain null, they have to move faster in the φ direction to cross the same temporal distance than nearby- ℓ geodesics (which move at finite speed in the ρ direction everywhere). Assuming that this effect dominates over what happens to the geodesics in passage between the boundary and vicinity of $\rho_0(\ell_0)$, the ℓ_0 geodesic will reach $\varphi = \pi$ ‘earlier’ than the nearby- ℓ geodesics. This in turn means that if we cut them off at $t = 0$ to find $\Xi_{\mathcal{A}}$, the pinch-off which happens at $\varphi = \pi$ will occur near the ℓ_0 geodesic. This implies that $\ell_* \approx \ell_0$. Assuming that in fact $\ell_* = \ell_0$, it is very simple to find the critical curve in $(r_h, \varphi_{\mathcal{A}})$ plane given by a function $\varphi_{\mathcal{A}}^*(r_h)$, since for each fixed r_h (which determines ℓ_0 using (2.5)), we merely need to find ρ_* by solving $t_{\ell_0}(\rho_*) = 0$, and then integrating the ℓ_0 -geodesic from $(t = 0, \rho = \rho_*, \varphi = \pi)$ back out to the boundary to find $\varphi_{\mathcal{A}}$.

We have used this trick to plot the critical curve on $(\varphi_{\mathcal{A}}, \rho_h)$ plane in figure 5 for various spacetime dimensions. The approximation $\ell_* \approx \ell_0$ can independently be checked by explicit numerical integration and it works extremely well for a large range of black hole sizes.¹⁵ As expected, for tiny black holes $\rho_h \ll 1$, the critical size of the region is $\varphi_{\mathcal{A}}^* \rightarrow \pi/2$ and grows linearly with $\rho_h \sim r_h$ since there is effectively no other scale in this regime, whereas for very large black holes $\rho_h \gg 1$, $\varphi_{\mathcal{A}}^* \rightarrow \pi$ asymptoting to a constant. Note that as $\varphi_{\mathcal{A}} \rightarrow \pi$ we are guaranteed to have a non-trivial topology of $\Xi_{\mathcal{A}}$. In particular, consider the limiting case where $\mathcal{A} = \mathbf{S}^3 \setminus i^0$, i.e., \mathcal{A} is a punctured sphere (henceforth denoted as $\varphi_{\mathcal{A}} = \pi^-$). Then no matter how large the black hole is, the causal wedge reaches all the way around the boundary while having a hole due to the horizon. However, from the observation of footnote 8 (further discussed in section 3) it follows that the limit $\varphi_{\mathcal{A}} \rightarrow \pi$ is not smooth.

It is worth remarking that the non-trivial topology of the causal wedge described above relies on working in global Schwarzschild-AdS geometry. In the planar Schwarzschild-AdS black hole geometry, one cannot circumnavigate the black hole, there being no “other side”. Nevertheless, as we will describe later, even in the Poincaré patch of AdS it is possible to encounter causal wedges with non-trivial topology; for instance a localized black hole in the Poincaré patch will likewise do the trick (see section 2.3).

Above we have presented an example of a causal wedge with one hole. It is now conceptually easy to generalize this situation to a causal wedge with multiple holes. For

¹⁵ In examples that we have examined we find that $\ell_* - \ell_0 \sim 10^{-3}$ for black holes which are roughly of the order of the AdS radius.




Figure 5. The critical curves on $(\varphi_{\mathcal{A}}, \rho_h)$ plane indicating where $\Xi_{\mathcal{A}}$ pinches off for Schwarzschild-AdS $_{d+1}$. $\Xi_{\mathcal{A}}$ has two components above the curve and only a single component below. To guide the eye, we also indicate the $\varphi_{\mathcal{A}} = \pi/2$ and $\varphi_{\mathcal{A}} = \pi$ (dashed lines); the latter gives the upper bound in $\varphi_{\mathcal{A}}$, while the former indicates the lower bound below which $\Xi_{\mathcal{A}}$ is connected for Schwarzschild-AdS black hole of any size. The topmost (red) curve corresponds to Schwarzschild-AdS $_4$ geometry where the effect of gravity is strongest, while the next (orange) curve is for Schwarzschild-AdS $_5$ which is our prime exhibit. Increasing the dimension results in slower growth of $\varphi_{\mathcal{A}}^*(\rho_h)$, as exemplified by $d = 5$ (green), $d = 6$ (blue), $d = 7$ (purple), $d = 19$ (purple dotted) and $d = 49$ (blue dashed). The effects of the weaker gravitational potential are clearly visible with the increasing dimension and the bottommost curve is close to the limiting behaviour for large d .

example, we can consider a (dynamical) situation with multiple small black holes in AdS.¹⁶ The black holes will generically orbit each other on timescales set by their separation, radiate gravitational waves, and eventually coalesce. But we can separate scales in such a way that around each black hole there is a region which is inside a causal wedge for appropriate \mathcal{A} . In fact, if $\varphi_{\mathcal{A}} \approx \pi/2$, we may see transitions in the number of components of Ξ . For small black holes and $\varphi_{\mathcal{A}}$ larger than $\pi/2$ by amount related to the black hole separation, $\Xi_{\mathcal{A}}$ will have a component around each horizon, apart from the one connected to $\partial\mathcal{A}$. Hence for a ‘galaxy’ with N separated black holes in AdS, Ξ will have $N + 1$ disconnected components.

So far, this section has focused on the topology of the causal wedge and the connectedness of $\Xi_{\mathcal{A}}$. Before closing, let us make an observation about the nature of the ‘phase transition’ between connected and disconnected $\Xi_{\mathcal{A}}$ as seen by its area $\chi_{\mathcal{A}}$. Although we don’t evaluate the causal holographic information χ explicitly, we expect that for fixed r_h ,

¹⁶ While such solutions are not known analytically one can construct approximate solutions using a matched asymptotics method.

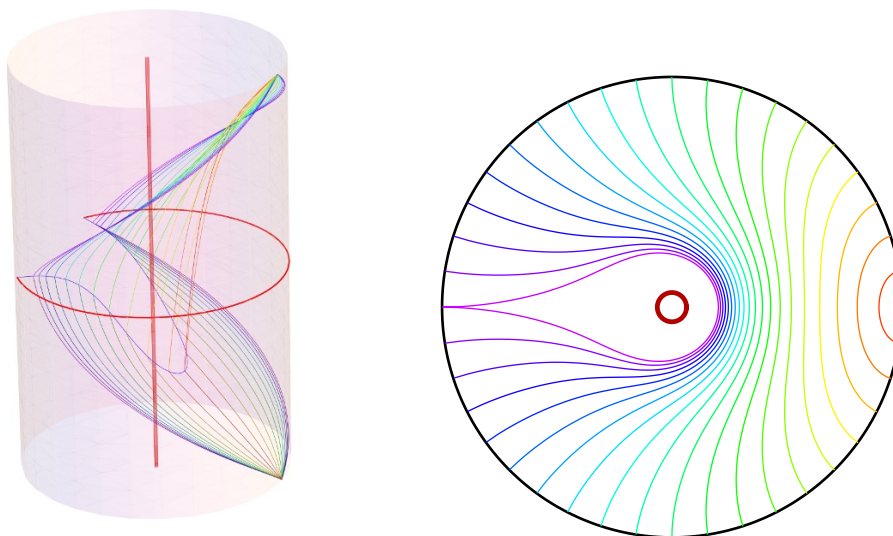


Figure 6. (Left:) A plot of the causal wedge $\diamond_{\mathcal{A}}$ and $\Xi_{\mathcal{A}}$ (blue curve) in BTZ, with $r_h = 0.02$ and $\varphi_{\mathcal{A}} = 2.5$. (Right:) Projection of spacelike geodesics $\mathfrak{E}_{\mathcal{A}}$ onto the Poincaré disk, for varying $\varphi_{\mathcal{A}} \in [0, \pi]$ in increments of 0.05π and $r_h = 0.1$. Since $\mathfrak{E}_{\mathcal{A}} = \Xi_{\mathcal{A}}$ in BTZ, the left and right panels are analogous to figure 3 and figure 4, respectively (modulo the different r_h values).

$\chi_{\mathcal{A}}(\varphi_{\mathcal{A}})$ is not smooth at the transition point $\varphi_{\mathcal{A}}^*$. Consider the full curve (on our (t, ρ, φ) subspace) $\mathcal{X}_{t=0}$ generated by intersections of future and past geodesics, characterized by $(\rho_{t=0}(\ell), \varphi_{t=0}(\ell))$. The causal information surface $\Xi_{\mathcal{A}}$ is a subset of this curve, restricted by $\varphi_{t=0}(\ell) \leq \pi$; in particular the two curves are identical only when $\Xi_{\mathcal{A}}$ has just one connected component. Since the spacetime is smooth, the geodesics, and hence their intersections, must vary smoothly in ℓ . Similarly, the length of the full $\mathcal{X}_{t=0}$ should vary smoothly as we change $\varphi_{\mathcal{A}}$. However, at the transition point $\varphi_{\mathcal{A}}^*$, χ ceases to be given by length of $\mathcal{X}_{t=0}$: the difference is given by the finite length piece where $\varphi_{t=0}(\ell) > \pi$. We would therefore expect that $\chi_{\mathcal{A}}(\varphi_{\mathcal{A}})$ has a kink at $\varphi_{\mathcal{A}} = \varphi_{\mathcal{A}}^*$.

2.2 BTZ

Having seen the rich structure of $\Xi_{\mathcal{A}}$ for Schwarzschild-AdS₅, one might wonder whether it is also present in the simpler case of the 3-dimensional BTZ geometry, which has the form (2.2) with

$$g(\rho) = \sin^2 \rho - r_h^2 \cos^2 \rho . \tag{2.11}$$

Here the calculation is in fact much simpler, and the explicit expressions for null geodesics which generate $\partial \diamond_{\mathcal{A}}$ were presented in e.g. [20]. A representative causal wedge is plotted in the left panel of figure 6. We choose $\varphi_{\mathcal{A}} = 2.5$ as in figure 3, and we use a tiny black hole $r_h = 0.02$ in order to emphasize the difference from the higher-dimensional case. In fact, as observed already in [13], in BTZ spacetime $\Xi_{\mathcal{A}}$ exactly coincides with $\mathfrak{E}_{\mathcal{A}}$; the latter corresponds to the spacelike geodesic anchored at $\pm\varphi_{\mathcal{A}}$ on the boundary. The spatial

projection of these is described by

$$\tan \rho = r_h \frac{\cosh(r_h \varphi_{\mathcal{A}})}{\sqrt{\cosh^2(r_h \varphi_{\mathcal{A}}) - \cosh^2(r_h \varphi)}} . \quad (2.12)$$

As can be easily seen, connected spacelike geodesics always exist for arbitrary $\varphi_{\mathcal{A}}$, as illustrated in the right panel of figure 6.

One reason why the causal wedge does not close off as in the higher dimensional case is that taking the black hole arbitrarily small $r_h \rightarrow 0$ does not approach pure AdS: the latter is achieved when $r_h^2 = -1$. Said differently, in 3 dimensions, the influence of the black hole does not fall off fast enough. Not only are the effects of a tiny BTZ black hole perceptible on AdS scale, but even near the boundary there is a qualitative difference between presence and absence of a black hole. A different way to see that 3-dimensional bulk is special is to note that we have argued that causal wedges must be simply connected to the boundary domain \mathcal{A} [19] as follows from topological censorship [23]. Since in this low dimension $\Xi_{\mathcal{A}}$ is a curve, it must be smoothly deformable to $\diamond_{\mathcal{A}}$. It thus follows that there is no room for non-trivial topology of $\Xi_{\mathcal{A}}$ in three dimensional bulk spacetimes.

2.3 Boosted black hole

In section 2.1 we have seen that for higher dimensional Schwarzschild-AdS black hole of any size, the causal wedge has holes for sufficiently large boundary region \mathcal{A} . As $r_h \rightarrow \infty$, the critical size of \mathcal{A} for which $\Xi_{\mathcal{A}}$ becomes disconnected approaches $\varphi_{\mathcal{A}}^* \rightarrow \pi$, whereas for $r_h \rightarrow 0$, the critical size $\varphi_{\mathcal{A}}^* \rightarrow \pi/2$. In particular, to obtain disconnected $\Xi_{\mathcal{A}}$ for this class of geometries, the region \mathcal{A} must cover at least half of the boundary sphere, and therefore sample a large part of the system. However, we now argue that large \mathcal{A} is actually not a prerequisite for existence of disconnected $\Xi_{\mathcal{A}}$, in the sense that for *any* finite region \mathcal{A} on the boundary we can construct asymptotically AdS geometries (in more than 3 dimensions) for which $\Xi_{\mathcal{A}}$ is disconnected.

In fact, a simple example which does the job is a boosted version of the global Schwarzschild-AdS black hole discussed in section 2.1. We can consider a family of geometries, considered e.g. in [1], corresponding to a boosted global black hole with fixed total energy.¹⁷ At zero boost this is the standard global AdS black hole while at infinite boost, this solution limits to a gravitational shock wave in AdS. In the static coordinates (defined with respect to a specified boundary time), the boosted black hole follows a trajectory which approximates that of a timelike geodesic.

We don't need to do a new calculation to see what will happen for causal wedges in such a geometry, since we can simply boost our causal wedge for the static black hole found in section 2.1. In other words, we can implement a coordinate transformation which in pure AdS would transform a timelike geodesic at the origin $\rho(t) = 0$ to one which oscillates back and forth with energy $E > 1$, whose radial profile is given by

$$\rho(t) = \sin^{-1} \left[\frac{\sqrt{E^2 - 1}}{E} \sin t \right] . \quad (2.13)$$

¹⁷ See also [30] for a recent discussion where such geometries were called oscillons.

The requisite transformation is most easily obtained from isometrically embedding AdS_{d+1} into $\mathbf{R}^{d,2}$ endowed with the flat metric $ds^2 = -dX_{-1}^2 - dX_0^2 + \sum_{i=1}^d dX_i^2$, restricted to the hyperboloid $-X_{-1}^2 - X_0^2 + \sum_{i=1}^d X_i^2 = -1$. This embedding makes the AdS isometries obvious: for instance we have manifest boost invariance, say $X_0 \rightarrow \cosh \beta X_0 + \sinh \beta X_1$ and $X_1 \rightarrow \cosh \beta X_1 + \sinh \beta X_0$, leaving all the other X_i 's unchanged. The energy E in (2.13) is related to the boost in the obvious manner $E = \cosh \beta$. So the relevant isometric embedding, which implements the boost and yields AdS in the conformally ESU coordinates, i.e. (2.2) with $g(\rho) = 1$, is

$$\begin{aligned} X_0 &= \cosh \beta \frac{\sin t}{\cos \rho} + \sinh \beta \frac{\sin \rho}{\cos \rho} \cos \varphi, & X_{-1} &= \frac{\cos t}{\cos \rho}, \\ X_1 &= \cosh \beta \frac{\sin \rho}{\cos \rho} \cos \varphi + \sinh \beta \frac{\sin t}{\cos \rho}, & X_k &= \frac{\sin \rho}{\cos \rho} \sin \varphi \Omega_k, \end{aligned} \quad (2.14)$$

where Ω_i with $i = 2, \dots, d$ are direction cosines, i.e., $\sum_i \Omega_i^2 = 1$, coordinatizing a unit \mathbf{S}^{d-2} and thus explicitly ensuring the $\text{SO}(d-1)$ symmetry in the transverse space. The actual transformation is then generated by comparing the X 's for $\beta = 0$ with those for arbitrary boost in (2.14), and is given (modulo some branch issues) by

$$\begin{aligned} \bar{\rho}(\rho, t, \varphi) &= \tan^{-1} \left[\frac{1}{\cos \rho} \sqrt{(\cosh \beta \sin \rho \cos \varphi + \sinh \beta \sin t)^2 + \sin^2 \rho \sin^2 \varphi} \right] \\ \bar{t}(\rho, t, \varphi) &= \tan^{-1} \left[\cosh \beta \tan t + \sinh \beta \frac{\sin \rho \cos \varphi}{\cos t} \right] \\ \bar{\varphi}(\rho, t, \varphi) &= \cot^{-1} \left[\cosh \beta \cot \varphi + \sinh \beta \frac{\sin t}{\sin \rho \sin \varphi} \right] \end{aligned} \quad (2.15)$$

where the barred coordinates correspond to boost with respect to the unbarred coordinates by a boost β . The direction cosines are unchanged since we retain $\text{SO}(d-1)$ symmetry.

Note that the boundary is preserved, $\bar{\rho}(\rho = \frac{\pi}{2}, t, \varphi) = \frac{\pi}{2}$ under (2.15), and this then specifies the corresponding transformation induced on the boundary. More precisely, the bulk coordinate transform implementing the boost corresponds to a conformal transformation (involving both time and space) on the boundary. This transformation changes the size of $\diamond_{\mathcal{A}}$. In particular, small \mathcal{A} and highly boosted black hole translates to large \mathcal{A} in the static black hole frame, at time where the black hole is closest to the region \mathcal{A} on the boundary in the boosted picture. This is illustrated in figure 7, which shows how a fixed size region \mathcal{A} taken at different times translates into a variable-sized region in a boosted frame with relative boost β (red curve corresponds to zero boost while purple to $\beta = 4$). If we take a hemispherical region with $\bar{\varphi}_{\mathcal{A}} = \pi/2$ (left), the ‘breathing’ of its size $\varphi_{\mathcal{A}}$ in the boosted frame is symmetric, whereas starting with a small region (right) only produces large sizes in the boosted frame at short intervals. Nevertheless, it is easy to see that no matter how small we take the region \mathcal{A} in the unboosted frame, its size in the boosted frame can get arbitrarily close to π for sufficiently large boosts.

It is now a simple matter to translate the effect of the boost on the causal wedge: it is precisely the same effect as that of varying the size of $\varphi_{\mathcal{A}}$. The topological attributes cannot change by the coordinate transformation; so since $\diamond_{\mathcal{A}}$ has a hole in unboosted




Figure 7. Effect of boost on the size $\varphi_{\mathcal{A}}$ of a region \mathcal{A} , as a function of time, as described in the text. Specifically, we take a fixed-sized region $\bar{\varphi}_{\mathcal{A}} = \frac{\pi}{2}$ (Left) and $\bar{\varphi}_{\mathcal{A}} = \frac{\pi}{8}$ (Right) in the boosted frame and plot how its size $\varphi_{\mathcal{A}}$ varies in the unboosted frame as a function of time t , for various values of boost, from unboosted case $\beta = 0$ (red) to $\beta = 4$ (purple) in increments of 0.5. This illustrates how arbitrarily small region in boosted frame can look large in unboosted frame for sufficiently large boosts, albeit for time intervals whose size also shrinks with the boost.

frame (equivalently the static black hole spacetime) for large regions, the same is true in the boosted frame (equivalently in the boosted black hole spacetime) even for small regions — but only at the auspicious times and for sufficient boosts.

It is worth noting that another construction which describes the same (Schwarzschild-AdS) geometry in different coordinates is obtained by simply restricting attention to the Poincaré patch of global Schwarzschild-AdS. The CFT dual on Minkowski space of such configurations was dubbed ‘the conformal soliton flow’ by [24]. The event horizon defined with respect to the Poincaré patch is then closely analogous to the boundary of the causal wedge, since it is generated by null geodesics in Schwarzschild-AdS which end on the boundary at i^+ of the Poincaré patch which corresponds to some finite ESU time; this was constructed explicitly in [31] (cf., their figure 3 for a plot¹⁸ of the event horizon in global coordinates in 3 dimensions, though the vertical axis is the BTZ time t rather than the coordinate used above in figure 6 which is more analogous to Eddington time). More precisely, the future event horizon indicated there would correspond to $\partial_+(\diamond_{\mathcal{A}})$ for $\varphi_{\mathcal{A}} = \pi$, though $\Xi_{\mathcal{A}}$ for any other $\varphi_{\mathcal{A}}$ can easily be read off by slicing their blue surface horizontally higher than i^0 .

So far in this subsection we have discussed geometries which are static though not manifestly so (they all admit a Killing field which is timelike everywhere outside of event horizon). It is even easier to construct examples with disconnected $\Xi_{\mathcal{A}}$ for dynamical situations with collapsing and evaporating black holes, where the event horizon lasts only for a finite time. Again by separation of scales, if the black holes are tiny on AdS scale we can put them anywhere, and they will only eat out tiny spacetime regions from the AdS causal wedge. Then the original $\partial\diamond_{\mathcal{A}}$ (connected to $\partial\lozenge_{\mathcal{A}}$) will be only minimally deformed

¹⁸ Strictly speaking the plot in [31] is for the BTZ conformal soliton where we have argued for the absence of non-trivial topology; a similar picture in higher dimensions should reveal the appropriate structure.

but the full $\partial\blacklozenge_{\mathcal{A}}$ will now include additional disconnected components. This will again translate into $\Xi_{\mathcal{A}}$ having multiple components, but now if we try to project them onto a single Poincaré disk, they may even intersect. However, topological censorship will still guarantee that the causal wedge remains simply connected.

Finally, note that since we can take \mathcal{A} arbitrarily small, the above construction can likewise be implemented in Poincaré AdS, where the dual CFT lives on flat space. In other words, having asymptotically global AdS bulk geometry is not a prerequisite to disconnected $\Xi_{\mathcal{A}}$ either.

2.4 Stars in AdS

So far all our examples of disconnected $\Xi_{\mathcal{A}}$ involved causally non-trivial bulk geometries. In such situations it is easy to argue from first principles that for sufficiently large \mathcal{A} the causal wedge must have holes, since by causality $\blacklozenge_{\mathcal{A}}$ cannot reach past the event horizon. However, we will now see that the presence of an event horizon is not a prerequisite for disconnected $\Xi_{\mathcal{A}}$. As long as $\mathcal{A}^c \neq \emptyset$ and we are in more than 3 bulk dimensions, all of the examples discussed above can be modified to keep the geometry causally trivial but nevertheless the causal wedge unaffected. In black hole geometries, null geodesics starting arbitrarily near the event horizon take a long time to reach the AdS boundary. So they typically don't make it out to $\blacklozenge_{\mathcal{A}}$; as a result $\Xi_{\mathcal{A}}$ remains a finite distance away from the event horizon.¹⁹ We could then take the metric to remain identical in the region reached by $\blacklozenge_{\mathcal{A}}$ but modify it outside that region so as to get rid of all the horizons.

For example in the Schwarzschild-AdS case, we could in principle replace the black hole by a compact star (or even a static shell) which is just slightly bigger than the original black hole. To find out whether or not this is physically realistic, let us examine how compact would such a star have to be. Since its maximal size is bounded by the deepest reach of $\Xi_{\mathcal{A}}$, and the latter occurs at $\varphi = 0$, we merely need to see how deep does the radial ($\ell = 0$) null geodesic from q^\vee penetrate by the time $t = 0$. This value of course depends on both $\varphi_{\mathcal{A}}$ and r_h . When $\varphi_{\mathcal{A}}$ is small, the geodesic of course does not have time to reach very deep, but then the causal wedge does not have holes, so this regime is irrelevant for our purposes. On the other hand, if we take $\varphi_{\mathcal{A}}$ too large in order to guarantee $\blacklozenge_{\mathcal{A}}$ having a hole, the geodesics have longer time to travel and typically they approach exponentially close to the horizon. For example in the case illustrated in figure 3, the ρ value reached by the radial (red) geodesics ρ_{Ξ} is already very close to the horizon: $\frac{\rho_{\Xi} - \rho_h}{\rho_h} \approx 5 \times 10^{-6}$, which is scarcely realistic for a compact star.

To maximize the size of the requisite star, then, we want to minimize $\varphi_{\mathcal{A}}$ subject to $\Xi_{\mathcal{A}}$ being disconnected, i.e. take $\varphi_{\mathcal{A}} = \varphi_{\mathcal{A}}^*$, and study the corresponding ρ_{Ξ} as a function of ρ_h . Numerical studies indicate that the maximal value attained, which occurs in the limit of tiny black holes, is $\frac{\rho_{\Xi} - \rho_h}{\rho_h} \approx 10^{-4}$. While this is still not be achievable for physically relevant equations of state for the star, it demonstrates the matter-of-principle point that event horizons are not necessary for topologically non-trivial causal wedge. In fact, preliminary

¹⁹ This is an extreme case of the time delay effect discussed in [3].

investigations indicate that it might be possible for charged scalar solitons in AdS to have non-trivial causal wedge topology; it would be interesting to explore this further.

3 General features of $\Xi_{\mathcal{A}}$ and $\chi_{\mathcal{A}}$

Having seen some curious features of causal wedges in specific classes of examples, we now turn to explaining the more ‘standard’ and completely general properties indicated in section 1.2. We consider generic boundary regions \mathcal{A} and \mathcal{B} (using the subscript t to indicate when they lie on the same time slice) and describe simple relational properties²⁰ of causal wedges etc. associated with them, including observations pertaining to our constructs in causally non-trivial spacetimes. We will then compare the areas of various surfaces and consider (sub)additivity properties of χ . Finally, we will close with a discussion of extremal surfaces, both within the boundary of the causal wedge and in the full spacetime. In the process, we will specify a useful relation between $\diamond_{\mathcal{A}}$ and any extremal surface $\mathfrak{E}_{\mathcal{A}}$, whose implications we will consider in section 4. In order to facilitate the reading, in section 3.1 we summarize our claims, leaving their proofs and discussion to section 3.2.

3.1 Summary of properties

The simple properties which we prove in section 3.2 are as follows:

1. If $\mathcal{A}_t \cap \mathcal{B}_t = \emptyset$ (and more generally if \mathcal{A} and \mathcal{B} are spacelike-separated), then $\diamond_{\mathcal{A}}$ and $\diamond_{\mathcal{B}}$ are spacelike-separated. Hence $\diamond_{\mathcal{A}} \cap \diamond_{\mathcal{B}} = \emptyset$ and $\Xi_{\mathcal{A}}$ and $\Xi_{\mathcal{B}}$ are likewise spacelike-separated.
2. If $\mathcal{A}_t \subset \mathcal{B}_t$ (and more generally if $\diamond_{\mathcal{A}} \subset \diamond_{\mathcal{B}}$), then $\diamond_{\mathcal{A}} \subset \diamond_{\mathcal{B}}$. Moreover, if \mathcal{A}_t is entirely inside \mathcal{B}_t (more generally if $\partial \diamond_{\mathcal{A}} \cap \partial \diamond_{\mathcal{B}} = \emptyset$), then $\Xi_{\mathcal{A}}$ and $\Xi_{\mathcal{B}}$ are spacelike-separated, with $\Xi_{\mathcal{B}}$ lying deeper than $\Xi_{\mathcal{A}}$ (equivalently outside $\diamond_{\mathcal{A}}$).
3. If \mathcal{A}_t and \mathcal{B}_t overlap (i.e., if $\mathcal{A} \cap \mathcal{B} \neq \emptyset$, but $\mathcal{A} \setminus \mathcal{B} \neq \emptyset$ and $\mathcal{B} \setminus \mathcal{A} \neq \emptyset$), then

$$\diamond_{\mathcal{A} \cap \mathcal{B}} \subset \diamond_{\mathcal{A}} \cap \diamond_{\mathcal{B}} \subset \{\diamond_{\mathcal{A}}, \diamond_{\mathcal{B}}\} \subset \diamond_{\mathcal{A}} \cup \diamond_{\mathcal{B}} \subset \diamond_{\mathcal{A} \cup \mathcal{B}}. \quad (3.1)$$

where by $\{\diamond_{\mathcal{A}}, \diamond_{\mathcal{B}}\}$ we mean either $\diamond_{\mathcal{A}}$ or $\diamond_{\mathcal{B}}$. Moreover, $\Xi_{\mathcal{A} \cup \mathcal{B}}$ and $\Xi_{\mathcal{A} \cap \mathcal{B}}$ are spacelike-separated, with $\Xi_{\mathcal{A} \cup \mathcal{B}}$ lying deeper than $\Xi_{\mathcal{A} \cap \mathcal{B}}$ (i.e. outside $\diamond_{\mathcal{A} \cap \mathcal{B}}$), etc.

4. For any (not necessarily stationary) bulk black hole spacetimes, no causal wedge can penetrate the event horizon. It then immediately follows that $\Xi_{\mathcal{A}}$ cannot penetrate an event horizon for any \mathcal{A} .
5. Causal holographic information satisfies Additivity:
If $\mathcal{A} \cap \mathcal{B} = \emptyset$, then $\chi_{\mathcal{A} \cup \mathcal{B}} = \chi_{\mathcal{A}} + \chi_{\mathcal{B}}$.
6. Causal holographic information satisfies Subadditivity:
If $\mathcal{A} \cap \mathcal{B} \neq \emptyset$, then $\chi_{\mathcal{A} \cup \mathcal{B}} \leq \chi_{\mathcal{A}} + \chi_{\mathcal{B}}$.
(However, Strong Subadditivity is *not* necessarily satisfied, as demonstrated in [13].)

²⁰ See also e.g. [3, 18, 19] for earlier related discussions.

- 7. From the set of all surfaces on $\partial\blacklozenge_{\mathcal{A}}$ anchored on $\partial\mathcal{A}$, $\Xi_{\mathcal{A}}$ is the minimal-area one.
- 8. Extremal surface $\mathfrak{E}_{\mathcal{A}}$ must lie outside (or on the boundary of) the causal wedge $\blacklozenge_{\mathcal{A}}$.²¹

3.2 Proofs of simple properties

Some of the proofs of the above statements use the concept of a causal curve, which we take to be a nowhere-spacelike, maximally extended, connected curve (either in the bulk or along the boundary). We will denote a future-directed causal curve from a point p to a point q in the spacetime by $\gamma_{p \rightarrow q}^+ = \gamma_{q \rightarrow p}^-$. Existence of such a curve guarantees that q is in the future of p .

(1). Causal wedges of spacelike-separated regions are spacelike-separated: First note that if \mathcal{A} and \mathcal{B} are spacelike-separated, then the corresponding domains of dependence $\blacklozenge_{\mathcal{A}}$ and $\blacklozenge_{\mathcal{B}}$ are spacelike-separated. Otherwise there would exist a (boundary) causal curve which contains points in $\blacklozenge_{\mathcal{A}}$ (and hence, by definition of domain of dependence, must intersect \mathcal{A}) as well as points in $\blacklozenge_{\mathcal{B}}$ (and therefore must intersect \mathcal{B} as well). However, existence of causal curve through both \mathcal{A} and \mathcal{B} contradicts the assumption that \mathcal{A} and \mathcal{B} are spacelike-separated.

We can now extend essentially the same proof-by-contradiction into the bulk: If $\blacklozenge_{\mathcal{A}}$ and $\blacklozenge_{\mathcal{B}}$ are not spacelike-separated, then there exists a (bulk) causal curve $\gamma_{a \rightarrow b}$ passing through points $a \in \blacklozenge_{\mathcal{A}}$ and $b \in \blacklozenge_{\mathcal{B}}$. w.l.o.g. assume that $\gamma_{a \rightarrow b} = \gamma_{a \rightarrow b}^+$ is future-directed. By definition of causal wedge, we also know that there exists a causal curve $\tilde{\gamma}_a$ through the point a which starts and ends on the boundary inside $\blacklozenge_{\mathcal{A}}$ and similarly $\exists \tilde{\gamma}_b$ through b starting and ending in $\blacklozenge_{\mathcal{B}}$. Out of these three causal curves we can now create a new causal curve $\gamma = \tilde{\gamma}_{\blacklozenge_{\mathcal{A}} \rightarrow a}^+ \cup \gamma_{a \rightarrow b}^+ \cup \tilde{\gamma}_{b \rightarrow \blacklozenge_{\mathcal{B}}}^+$ composed of the past part of $\tilde{\gamma}_a$, $\gamma_{a \rightarrow b}^+$, and future part of $\tilde{\gamma}_b$, joined at a and b , starting from $\blacklozenge_{\mathcal{A}}$ and ending in $\blacklozenge_{\mathcal{B}}$. By suitably projecting γ onto the boundary obtains a causal²² curve $\bar{\gamma}$ intersecting both $\blacklozenge_{\mathcal{A}}$ and $\blacklozenge_{\mathcal{B}}$, contradicting the observation that $\blacklozenge_{\mathcal{A}}$ and $\blacklozenge_{\mathcal{B}}$ are spacelike-separated. This proves that $\blacklozenge_{\mathcal{A}}$ and $\blacklozenge_{\mathcal{B}}$ are spacelike-separated.

In particular, there are no points in $\blacklozenge_{\mathcal{A}}$ and $\blacklozenge_{\mathcal{B}}$ which can be connected by a causal curve. Since the causal information surfaces Ξ are contained in the causal wedges, there is correspondingly no causal curve connecting $\Xi_{\mathcal{A}}$ and $\Xi_{\mathcal{B}}$ — hence these are likewise spacelike-separated. Moreover, the absence of causal curve connecting points in $\blacklozenge_{\mathcal{A}}$ and $\blacklozenge_{\mathcal{B}}$ trivially implies the absence of common points between $\blacklozenge_{\mathcal{A}}$ and $\blacklozenge_{\mathcal{B}}$; in other words, $\blacklozenge_{\mathcal{A}} \cap \blacklozenge_{\mathcal{B}} = \emptyset$. \square

(2). Causal wedge inclusion for nested regions: First of all, note that if $\mathcal{A} \subset \mathcal{B}$, then any boundary causal curve $\bar{\gamma}_a$ through a point $a \in \blacklozenge_{\mathcal{A}}$ must by definition intersect

²¹ However, as demonstrated in [20], $\mathfrak{E}_{\mathcal{A}}$ and $\Xi_{\mathcal{A}}$ need not be always spacelike-separated. When they are, the above statement guarantees that $\mathfrak{E}_{\mathcal{A}}$ lies deeper than $\Xi_{\mathcal{A}}$.

²² The fact that bulk causal curves ‘project’ to boundary causal curves is easy to see: the tangent vector to the bulk causal curve is timelike ($ds^2 < 0$) and in restricting to the boundary one eliminates a spatial direction which (as long as the projection is performed in such a way as to maintain the relative weighing of temporal and angular components) makes ds^2 even more negative. Also note that we are restricting to the causal wedge of the boundary and so are necessarily outside any black hole horizons.

\mathcal{A} and therefore it must necessarily also intersect \mathcal{B} , which means that $a \in \diamond_{\mathcal{B}}$. In other words, $\diamond_{\mathcal{A}} \subset \diamond_{\mathcal{B}}$. We now extend the same argument into the bulk: for any point $a \in \diamond_{\mathcal{A}}$, there exists a causal curve γ_a which begins and ends in $\diamond_{\mathcal{A}}$ and therefore begins and ends in $\diamond_{\mathcal{B}}$. This implies that $a \in \diamond_{\mathcal{B}}$, proving the inclusion $\diamond_{\mathcal{A}} \subset \diamond_{\mathcal{B}}$.

To show the rest of the statement, pertaining to $\diamond_{\mathcal{A}}$ lying strictly inside $\diamond_{\mathcal{B}}$, we first recall an obvious characteristic of points lying on the boundary of a causal wedge, namely the existence of nearby points which lie outside the causal wedge. In particular, since $\Xi_{\mathcal{B}}$ lies on the boundary of $\diamond_{\mathcal{B}}$, if $p \in \Xi_{\mathcal{B}}$, then within any open neighborhood $\mathcal{O}(p)$, there exists a point q which does not lie within $\diamond_{\mathcal{B}}$. More specifically, since $\Xi_{\mathcal{B}}$ lies on the intersection of future and past boundaries of the causal wedge, there exist points $q \in \mathcal{O}(p)$ through which no causal curve γ_q can start or end in $\diamond_{\mathcal{B}}$. This in turn implies that through any $b \in \Xi_{\mathcal{B}}$, any causal curve γ_b can at best make it to $\partial \diamond_{\mathcal{B}}$, but not inside $\diamond_{\mathcal{B}} \setminus \partial \diamond_{\mathcal{B}}$.

We now show that $\Xi_{\mathcal{B}}$ lies outside $\diamond_{\mathcal{A}}$. Using the previous observation, we can see that if any point $b \in \Xi_{\mathcal{B}}$ lies inside $\diamond_{\mathcal{A}}$, then there exists a causal curve γ_b which reaches $\diamond_{\mathcal{A}} \subset \diamond_{\mathcal{B}} \setminus \partial \diamond_{\mathcal{B}}$, a contradiction. To say this differently, suppose that there exists a point $b \in \Xi_{\mathcal{B}}$ which lies inside $\diamond_{\mathcal{A}}$. If $b \in \diamond_{\mathcal{A}} \setminus \partial \diamond_{\mathcal{A}}$, then there exists an open neighborhood $\mathcal{O}(b)$ such that $\mathcal{O}(b) \subset \diamond_{\mathcal{A}} \subset \diamond_{\mathcal{B}}$, contradicting the assumption that $b \in \Xi_{\mathcal{B}}$. On the other hand, if no part of $\Xi_{\mathcal{B}}$ lies strictly inside $\diamond_{\mathcal{A}}$ but $b \in \partial \diamond_{\mathcal{A}}$, then $\Xi_{\mathcal{B}}$ must be tangent to $\partial \diamond_{\mathcal{A}}$. This means that the generator of $\partial \diamond_{\mathcal{B}}$ through b must coincide with the corresponding generator of $\partial \diamond_{\mathcal{A}}$. Since this generator must extend all the way to the boundary, it terminates on $\partial \diamond_{\mathcal{B}}$ and simultaneously on $\partial \diamond_{\mathcal{A}}$; but this contradicts our assumption that $\partial \diamond_{\mathcal{A}} \cap \partial \diamond_{\mathcal{B}} = \emptyset$. This argument also shows that the boundaries of the two causal wedges cannot coincide at any point, so $\diamond_{\mathcal{A}}$ must be strictly inside $\diamond_{\mathcal{B}}$.

Having proved that $\diamond_{\mathcal{A}}$ lies strictly inside $\diamond_{\mathcal{B}}$ and that $\Xi_{\mathcal{B}}$ lies outside $\diamond_{\mathcal{A}}$, the argument that $\Xi_{\mathcal{A}}$ and $\Xi_{\mathcal{B}}$ must be spacelike-separated proceeds analogously: w.l.o.g. suppose that there exists a future-directed causal curve $\gamma_{b \rightarrow a}^+$ (the proof for past-directed curves proceeds analogously) connecting a point $b \in \Xi_{\mathcal{B}}$ and a point $a \in \Xi_{\mathcal{A}}$. Then there is a causal curve $\gamma_b = \gamma_{b \rightarrow a}^+ + \tilde{\gamma}_{a \rightarrow \diamond_{\mathcal{A}}}^+$ through b which ends in $\diamond_{\mathcal{A}} \in \diamond_{\mathcal{B}} \setminus \partial \diamond_{\mathcal{B}}$, again a contradiction. Finally, since $\Xi_{\mathcal{A}}$ and $\Xi_{\mathcal{B}}$ are spacelike-separated and $\Xi_{\mathcal{B}}$ lies outside $\diamond_{\mathcal{A}}$, it immediately follows that $\Xi_{\mathcal{B}}$ lies deeper than $\Xi_{\mathcal{A}}$. \square

(3). Causal wedge inclusion for overlapping regions: Suppose that \mathcal{A} and \mathcal{B} overlap, so that $\mathcal{A} \cap \mathcal{B} \neq \emptyset$, with $\mathcal{A} \setminus \mathcal{B} \neq \emptyset$ and $\mathcal{B} \setminus \mathcal{A} \neq \emptyset$. Then $\mathcal{A} \cap \mathcal{B}$ is a proper subset of each of \mathcal{A} and \mathcal{B} which are in turn each a proper subset of $\mathcal{A} \cup \mathcal{B}$, and we can directly apply the results of property (2) discussed above. In particular, $\diamond_{\mathcal{A} \cap \mathcal{B}} \subset \{\diamond_{\mathcal{A}}, \diamond_{\mathcal{B}}\} \subset \diamond_{\mathcal{A} \cup \mathcal{B}}$, so that $\diamond_{\mathcal{A} \cap \mathcal{B}} \subset \{\diamond_{\mathcal{A}}, \diamond_{\mathcal{B}}\} \subset \diamond_{\mathcal{A} \cup \mathcal{B}}$, and moreover $\Xi_{\mathcal{A} \cap \mathcal{B}}$ and $\Xi_{\mathcal{A} \cup \mathcal{B}}$ are spacelike-separated, with $\Xi_{\mathcal{A} \cup \mathcal{B}}$ lying deeper than $\Xi_{\mathcal{A} \cap \mathcal{B}}$.

To see the ‘intermediate’ inclusions, it is evident that $\diamond_{\mathcal{A}} \cap \diamond_{\mathcal{B}} \subset \{\diamond_{\mathcal{A}}, \diamond_{\mathcal{B}}\} \subset \diamond_{\mathcal{A}} \cup \diamond_{\mathcal{B}}$, and all that remains to show is that $\diamond_{\mathcal{A} \cap \mathcal{B}} \subset \diamond_{\mathcal{A}} \cap \diamond_{\mathcal{B}}$ and that $\diamond_{\mathcal{A}} \cup \diamond_{\mathcal{B}} \subset \diamond_{\mathcal{A} \cup \mathcal{B}}$. The arguments are similar to the ones used above: if $p \in \diamond_{\mathcal{A} \cap \mathcal{B}}$, then there exists a causal curve γ_p which starts and ends in $\diamond_{\mathcal{A} \cap \mathcal{B}}$, so it necessarily also starts and ends in each of $\diamond_{\mathcal{A}}$ and $\diamond_{\mathcal{B}}$, and therefore p lies in both $\diamond_{\mathcal{A}}$ and $\diamond_{\mathcal{B}}$ — which implies that $p \in \diamond_{\mathcal{A}} \cap \diamond_{\mathcal{B}}$. The other inclusion is similarly manifest. \square

(4). Causal wedge cannot penetrate event horizon: This statement follows immediately from the definition of an event horizon: since there is no future-directed causal curve from inside the black hole which can make out it to the AdS boundary, no point inside the black hole can lie inside the causal wedge of any boundary region. Correspondingly, the causal information surface $\Xi_{\mathcal{A}}$ for any sub-region \mathcal{A} of the total boundary space must lie outside the event horizon.²³

In asymptotically global AdS spacetimes, there is a slight subtlety: as we take \mathcal{A} to cover the entire spatial section of the boundary, $\diamond_{\mathcal{A}}$ jumps discontinuously from having finite time-extent (given by the size of the boundary sphere) to having infinite time extent (and covering the entire boundary spacetime). In this special case the causal wedge $\blacklozenge_{\mathcal{A}}$ is an open set, its boundary coincides with the event horizon, and $\Xi_{\mathcal{A}}$ then lies along the event horizon bifurcation surface. On the other hand, for \mathcal{A} being a proper subset of the boundary Cauchy slice, the finiteness of $\diamond_{\mathcal{A}}$ implies that $\Xi_{\mathcal{A}}$ can reach only to within a finite (albeit quantitatively small) distance from the horizon; we saw an example in section 2.4.

Now that we have discussed the relational properties between the causal wedges and the causal information surfaces for two regions, let us briefly turn to the causal holographic information χ . Recall that this quantity is potentially the one most directly accessible from the field theory.

(5). Additivity: If $\mathcal{A} \cap \mathcal{B} = \emptyset$, then by property (1), $\blacklozenge_{\mathcal{A}} \cap \blacklozenge_{\mathcal{B}} = \emptyset$, so $\Xi_{\mathcal{A}}$ and $\Xi_{\mathcal{B}}$ are disjoint. Hence $\Xi_{\mathcal{A} \cup \mathcal{B}} = \Xi_{\mathcal{A}} \cup \Xi_{\mathcal{B}}$, and the individual areas then simply add up, $\chi_{\mathcal{A} \cup \mathcal{B}} = \chi_{\mathcal{A}} + \chi_{\mathcal{B}}$. Note that each of the terms is divergent, with l.h.s. inheriting all the divergences from r.h.s. .

(6). Subadditivity: If $\mathcal{A} \cap \mathcal{B} \neq \emptyset$, then $\mathcal{A} \cup \mathcal{B}$ is not simply composed of two disjoint regions as in property (5) above. Instead, the surface area of this region is strictly smaller than the sum of the individual surface areas, $\text{Area}(\partial[\mathcal{A} \cup \mathcal{B}]) < \text{Area}(\partial[\mathcal{A}]) + \text{Area}(\partial[\mathcal{B}])$, because the r.h.s. has additional contribution from $\text{Area}(\partial[\mathcal{A} \cap \mathcal{B}])$. Now, since $\chi_{\mathcal{A}}$ has its leading divergence proportional to $\text{Area}(\partial[\mathcal{A}])$ and similarly for other regions, the Subadditivity inequality $\chi_{\mathcal{A} \cup \mathcal{B}} \leq \chi_{\mathcal{A}} + \chi_{\mathcal{B}}$ is satisfied trivially — i.e. the r.h.s. has stronger divergence. On the other hand, as noted in section 3.1, the property of Strong Subadditivity, where the leading divergences cancel, is actually not necessarily satisfied.

Properties (5) and (6), along with the basic feature of being (quarter of) the proper area of a co-dimension two surface, make it tempting to compare the causal holographic information $\chi_{\mathcal{A}}$ with the entanglement entropy $S_{\mathcal{A}}$. The failure of $\chi_{\mathcal{A}}$ to satisfy strong subadditivity is the strongest evidence that it cannot correspond to a von Neumann entropy.

²³ As an aside, note that this property of remaining outside the black hole does not generically hold for extremal surfaces, due to the teleological nature of the event horizon: an extremal surface cannot be sensitive to its exact location in dynamically evolving spacetimes, and therefore can probe inside the black hole; explicit examples have already been seen in [20, 32] and will be further discussed in [33]. (In contrast, as pointed out in [17], in static spacetimes extremal surfaces don't penetrate the event horizon either. Nevertheless, any extremal surface $\mathfrak{E}_{\mathcal{A}}$ must lie outside a causal wedge $\blacklozenge_{\mathcal{A}}$ as we show in Property 8.)

Finally, let us close with two properties which have been observed previously, and deal with extremal surfaces.

(7). $\Xi_{\mathcal{A}}$ is a minimal surface on $\partial\Diamond_{\mathcal{A}}$: We can in fact view this property as an alternate definition of $\Xi_{\mathcal{A}}$, which is conceptually useful for considering the causal holographic information $\chi_{\mathcal{A}}$. However, to establish the result, we use the definition (1.2), that $\Xi_{\mathcal{A}}$ belongs to both $\partial_+(\Diamond_{\mathcal{A}})$ and $\partial_-(\Diamond_{\mathcal{A}})$. The proof assumes null energy condition and uses a crucial observation about null congruences which generate the causal wedge boundary: the null congruences $\partial_{\pm}(\Diamond_{\mathcal{A}})$ must have non-negative (and generically positive) expansion in the outgoing (towards the boundary) direction. For if they had negative expansion, Raychaudhuri equation would imply that they caustic before reaching the AdS boundary, in contradiction to them generating the boundary of a causal past/future of $\Diamond_{\mathcal{A}} \in \partial\mathcal{M}$.

Let us first restrict attention to the case where $\Diamond_{\mathcal{A}}$ is topologically trivial, i.e. $\Xi_{\mathcal{A}}$ is a single connected surface anchored on $\partial\mathcal{A}$. Consider any other surface $\Upsilon_{\mathcal{A}} \subseteq \partial\Diamond_{\mathcal{A}}$ which is anchored on $\partial\mathcal{A}$. We want to show that $\Upsilon_{\mathcal{A}}$ cannot have smaller area than $\Xi_{\mathcal{A}}$. We can obtain a subset²⁴ of $\Upsilon_{\mathcal{A}}$ from $\Xi_{\mathcal{A}}$ by flowing a certain distance λ along the null generators. Let us for the moment assume that $\Upsilon_{\mathcal{A}} \in \partial_+(\Diamond_{\mathcal{A}})$; then we can perform a constant rescaling of the affine parameter of each null generator individually, such that $\Upsilon_{\mathcal{A}}$ lies at constant affine parameter λ_0 along the null generators of $\partial_+(\Diamond_{\mathcal{A}})$. Now, using the fact that the expansion of the null generators of $\partial_+(\Diamond_{\mathcal{A}})$ cannot be negative towards the boundary, we know that the area of constant λ slices of $\partial_+(\Diamond_{\mathcal{A}})$ must be monotonically increasing function of λ ; in particular, $\text{Area}(\Upsilon_{\mathcal{A}}) \geq \text{Area}(\Xi_{\mathcal{A}})$. (The presence of caustics in $\partial_+(\Diamond_{\mathcal{A}})$ would only strengthen this inequality, since the area of the remaining part of $\Upsilon_{\mathcal{A}}$, which does not lie along generators from $\Xi_{\mathcal{A}}$, is positive.) Same argument would apply for $\Upsilon_{\mathcal{A}}$ lying on $\partial_-(\Diamond_{\mathcal{A}})$, as the past-directed null generators again expand towards the boundary. If $\Upsilon_{\mathcal{A}}$ lies partly on $\partial_+(\Diamond_{\mathcal{A}})$ and partly on $\partial_-(\Diamond_{\mathcal{A}})$, then we can separate $\Upsilon_{\mathcal{A}}$ into domains, separated by $\Upsilon_{\mathcal{A}} \cap \Xi_{\mathcal{A}}$, and run the argument for each domain separately. Hence in all cases, any surface $\Upsilon_{\mathcal{A}}$ cannot have smaller area than $\Xi_{\mathcal{A}}$, which means that $\Xi_{\mathcal{A}}$ is the minimal surface on $\partial\Diamond_{\mathcal{A}}$.

Let us now turn to the topologically non-trivial case, such as discussed in section 2.1, where $\Diamond_{\mathcal{A}}$ has a hole and $\Xi_{\mathcal{A}}$ consists of several disconnected surfaces. Then the statement of minimality of $\Xi_{\mathcal{A}}$ has to be made more precise, as clearly there are surfaces $\Upsilon_{\mathcal{A}} \subseteq \partial\Diamond_{\mathcal{A}}$ with $\partial\Upsilon_{\mathcal{A}} = \partial\mathcal{A}$ with smaller area than $\Xi_{\mathcal{A}}$ — as a trivial example, take just the connected part of $\Xi_{\mathcal{A}}$ anchored on $\partial\mathcal{A}$, or a small deformation thereof. However, if in addition to $\partial\Upsilon_{\mathcal{A}} = \partial\mathcal{A}$ we require that $\Xi_{\mathcal{A}}$ and $\Upsilon_{\mathcal{A}}$ are homotopic to each other within $\partial\Diamond_{\mathcal{A}}$, then our arguments above go through: any generator of $\partial\Diamond_{\mathcal{A}}$ which starts from $\Xi_{\mathcal{A}}$ must intersect $\Upsilon_{\mathcal{A}}$, where the positive expansion of the generators guarantees that the latter has larger area. \square

The fact that $\Xi_{\mathcal{A}}$ is a minimal (and therefore extremal) surface on $\partial\Diamond_{\mathcal{A}}$ however does not mean that it is an extremal surface in the full bulk spacetime, as emphasized previously in [13, 20]. The final property will show this explicitly, by providing a relational property between $\mathfrak{E}_{\mathcal{A}}$ and $\Diamond_{\mathcal{A}}$.

²⁴ Recall that new generators can enter $\partial\Diamond_{\mathcal{A}}$ at caustics, so if $\Upsilon_{\mathcal{A}}$ lies closer to the boundary than the first caustic, all generators from $\Xi_{\mathcal{A}}$ pass through $\Upsilon_{\mathcal{A}}$ but in addition some new ones do as well.

(8). Extremal surface is generically outside causal wedge: This is a generalization of the proof presented in [13] (see also [34]), which was formulated on a given spacelike slice of the bulk rather than the full Lorentzian geometry.²⁵ The reasoning here is similar; the crux is to argue that if the extremal surface came within the causal wedge, it would have to be tangent to the boundary of a nested causal wedge corresponding to some sub-region. Once this is established, we can obtain a contradiction from comparing the expansions for the extremal surface and a tangent slice of a causal wedge boundary.

Suppose that an extremal surface $\mathfrak{E}_{\mathcal{A}}$ has a sub-region e which lies strictly inside $\diamond_{\mathcal{A}}$. We claim that then somewhere along e , there exists a point p at which $\mathfrak{E}_{\mathcal{A}}$ is tangent to $\partial\diamond_{\mathcal{B}}$ corresponding to some smaller boundary region $\mathcal{B} \subseteq \mathcal{A}$. There may in fact be infinitely many of such points, but our proof only requires one. The existence of one follows easily if one can foliate $\diamond_{\mathcal{A}}$ by $\partial\diamond_{\alpha}$ for a family of nested regions \mathcal{A}_{α} , parameterized by some parameter α , which we can choose to be $\alpha = 0$ at $\partial\diamond_{\mathcal{A}}$ and monotonically growing as the size of \mathcal{A}_{α} decreases.²⁶ For then this foliation specifies a function α on $\mathfrak{E}_{\mathcal{A}}$ (given by the particular $\partial\diamond_{\alpha}$ intersected by our surface $\mathfrak{E}_{\mathcal{A}}$ at any specified point). Since $\alpha = 0$ at $\mathfrak{E}_{\mathcal{A}} \cap \partial\diamond_{\mathcal{A}}$, i.e. on ∂e , and positive inside e , there is a point $p \in e$ where α takes extremal value. At this point p , $\mathfrak{E}_{\mathcal{A}}$ is tangent to $\partial\diamond_{\alpha}$.²⁷

Since the above argument followed from the existence of a foliation of $\diamond_{\mathcal{A}}$ by $\partial\diamond_{\alpha}$, we now pause to briefly consider how generally does such a foliation exist. First of all, it is not hard to see that for static spacetimes, a foliation of $\diamond_{\mathcal{A}}$ is indeed guaranteed to exist; as a canonical example (which includes caustics and potentially non-trivial topology of $\diamond_{\mathcal{A}}$), consider the Schwarzschild-AdS example discussed in section 2.1. Since the projections of $\Xi_{\mathcal{A}}(\varphi_{\mathcal{A}})$ on the Poincaré disk for different $\varphi_{\mathcal{A}}$ indicated in figure 4 foliate the entire region of the Poincaré disk which is in the largest causal wedge, this spatial foliation lifts trivially to the full Lorentzian region inside the largest causal wedge.²⁸ More generally, we can argue that such a foliation must exist for any static spacetime as follows. Let us w.l.o.g. fix \mathcal{A} to lie at $t = 0$ on the boundary. We first pick a convenient foliation of the region \mathcal{A} by slicing $\partial\diamond_{\mathcal{A}}$ by constant $t = \alpha$ surfaces and projecting to the $t = 0$ slice. Denote each

²⁵ In that more limited context, one can equivalently say that the extremal surface $\mathfrak{E}_{\mathcal{A}}$ reaches deeper than the causal information surface $\Xi_{\mathcal{A}}$; however, in general highly dynamical spacetimes such statement need not be meaningful when $\mathfrak{E}_{\mathcal{A}}$ and $\Xi_{\mathcal{A}}$ are not spacelike-separated.

²⁶ For example, one convenient parameterization is by the depth to which Ξ_{α} reaches, $\alpha = r_{\Xi_{\alpha}} - r_{\Xi_{\mathcal{A}}}$. Alternately, for spherically symmetric spacetimes such as considered in section 2.1, we can simply take $\alpha = \varphi_{\mathcal{A}} - \varphi_{\alpha}$.

²⁷A potential loophole is that $\partial\diamond_{\alpha}$ might be non-smooth at the point of intersection with $\mathfrak{E}_{\mathcal{A}}$, e.g. p could be one of the cusps seen in figure 4. In such cases there are two potential ways to extend our arguments. One is that we could consider other foliations of \mathcal{A} on the boundary to ensure that the intersection of the extremal surface with the causal wedge boundary is at a regular point. More simply however, we could also note that even at the cusps the expansion of the null congruence towards the boundary while divergent is manifestly positive definite (and we only need the correct sign for our argument to go through). We thank Matt Headrick and Aron Wall for useful discussions on this point.

²⁸Here we envision that $\diamond_{\mathcal{A}}$ is compact on the boundary. The issue is more subtle if we take \mathcal{A} to be a complete Cauchy slice of the boundary ESU; as noted in footnote 8, $\diamond_{\mathcal{A}}$ is the entire boundary spacetime and its causal wedge does not admit such a foliation (as apparent from both figure 4 and figure 6 (right panel)).

leaf of \mathcal{A} 's foliation by α and the enclosed subregions by \mathcal{A}_α . In case of circular regions, $\alpha \in (0, \varphi_{\mathcal{A}})$. Notice that the causal wedge \diamond_α for a given sub-region \mathcal{A}_α can be obtained by simply rigidly sliding $\partial_+(\diamond_{\mathcal{A}})$ and $\partial_-(\diamond_{\mathcal{A}})$ towards each other in time. Now consider some point $p \in \diamond_{\mathcal{A}}$. Since there is a unique point in $\partial_\pm(\diamond_{\mathcal{A}})$ whose temporal projection coincides with that of p , $\exists! \alpha$ by which we can slide $\partial_\pm(\diamond_{\mathcal{A}})$ such that $p \in \partial \diamond_\alpha$. In other words, this rigid time-translation of $\partial_\pm(\diamond_{\mathcal{A}})$ towards each other defines a natural foliation of $\diamond_{\mathcal{A}}$. We believe that it is possible to extend this proof to more general spacetimes, although we leave a rigorous argument for future investigation.

Let us now return to the main argument, having established that if a region $e \subset \mathfrak{E}_{\mathcal{A}}$ lies within $\diamond_{\mathcal{A}}$, then there exists a point $p \in e$ at which $\mathfrak{E}_{\mathcal{A}}$ is tangent to $\partial \diamond_{\mathcal{B}}$ for some sub-region \mathcal{B} . In fact, we can construct a spacelike surface $\Psi_{\mathcal{B}}$ within $\partial \diamond_{\mathcal{B}}$ which is tangent to $\mathfrak{E}_{\mathcal{A}}$ at p and is anchored on $\partial \mathcal{B}$. At this point, outgoing null normal to $\mathfrak{E}_{\mathcal{A}}$ coincides with the outgoing null normal to $\Psi_{\mathcal{B}}$ (which is the corresponding generator of $\partial \diamond_{\mathcal{B}}$). Now consider the expansions $\Theta_{\mathfrak{E}}$ and Θ_{Ψ} of the two tangent surfaces $\mathfrak{E}_{\mathcal{A}}$ and $\Psi_{\mathcal{B}}$ at p . By definition of extremal surface, we know $\Theta_{\mathfrak{E}} = 0$. On the other hand, since $\partial \mathcal{B}$ is a boundary of a causal set, $\Theta_{\Psi} \geq 0$ towards the boundary. This is however a contradiction, since the way in which these surfaces are tangent to each other (with \mathfrak{E} bending away from the boundary more than Ψ) implies that $\Theta_{\mathfrak{E}} > \Theta_{\Psi}$. The reader is encouraged to consult figure 5 of [13] for a pictorial sketch of this argument.

Since the assumption that $\mathfrak{E}_{\mathcal{A}}$ reaches inside $\diamond_{\mathcal{A}}$ produced a contradiction, we conclude that $\mathfrak{E}_{\mathcal{A}}$ must lie outside (or at best on the boundary of) $\diamond_{\mathcal{A}}$. \square

In the generic situation where $\Xi_{\mathcal{A}}$ is spacelike-separated from $\mathfrak{E}_{\mathcal{A}}$, it then immediately follows that $\mathfrak{E}_{\mathcal{A}}$ reaches deeper than $\Xi_{\mathcal{A}}$. However, having established that an extremal surface $\mathfrak{E}_{\mathcal{A}}$ cannot lie within the causal wedge, we should note that this does not automatically imply that $\mathfrak{E}_{\mathcal{A}}$ is necessarily causally disconnected from $\diamond_{\mathcal{A}}$ (although this is the case in generic situations, when $\mathfrak{E}_{\mathcal{A}}$ and $\Xi_{\mathcal{A}}$ are spacelike separated). Apart from the obvious special examples where $\mathfrak{E}_{\mathcal{A}}$ and $\Xi_{\mathcal{A}}$ coincide, as pointed out in footnote ²¹, in the case of thin shell Vaidya-AdS explored in [20] we saw that they can be null-separated.

4 Implications for bulk extremal surfaces

So far, we have been discussing the causal wedge and related constructs in the bulk space-time, which as yet have no independently defined construction in the dual CFT. In this section we point out that our results nevertheless bear on a more familiar context where we *do* have a conjectured duality. In particular, they surprisingly turn out to be relevant for the entanglement entropy $S_{\mathcal{A}}$, which is conjectured [15, 16, 21] to be given by (quarter of) the area of the extremal surface $\mathfrak{E}_{\mathcal{A}}$ which (i) is anchored on the entangling surface, $\partial \mathfrak{E}_{\mathcal{A}} = \partial \mathcal{A}$, (ii) is homologous to \mathcal{A} , and (iii) in case of multiple such surfaces is the minimal-area one.

In the previous section, we have further justified the anticipated result that no extremal surface $\mathfrak{E}_{\mathcal{A}}$ can lie within the causal wedge $\diamond_{\mathcal{A}}$ (Property 8). While we have presented this as a property of causal wedges, we can conversely think of it as a property of extremal surfaces. As already mentioned in section 1 the consequence that extremal surfaces penetrate deeper

into the bulk than causal wedges bears on the question of how much of the bulk does a given boundary region in the CFT describe. But quite apart from this discussion, there is a more remarkable and surprising consequence of Property 8, when combined with the observation of section 2 that causal wedges can have holes: in a global eternal black hole spacetime, if the causal wedge $\blacklozenge_{\mathcal{A}}$ for a given boundary region \mathcal{A} has a hole, then there *cannot* exist a connected extremal surface $\mathfrak{E}_{\mathcal{A}}$ anchored on $\partial\mathcal{A}$ which is homologous to \mathcal{A} .

To see why this is the case, let us first consider the homology requirement. Suppose we take a region \mathcal{A} which covers more than half of the boundary, such as the case indicated in figure 3. If area minimization was the only constraint on the desired extremal surface $\mathfrak{E}_{\mathcal{A}}$ anchored on $\partial\mathcal{A}$, then the extremal surface passing around the opposite side of the black hole from \mathcal{A} (i.e. through $\varphi = \pi$) would be the relevant one; let us denote it as $\mathfrak{E}_{\mathcal{A}^c}$. However, such a surface is not homologous to \mathcal{A} , since there does not exist a co-dimension one bulk smooth hypersurface whose only boundary are $\mathfrak{E}_{\mathcal{A}}$ and \mathcal{A} — such hypersurface would have to pass through the black hole, and would either encounter the black hole curvature singularity, or pass through the Einstein-Rosen bridge, in which case it would have further boundaries. The upshot is that, to satisfy the homology requirement, we should either take a surface $\mathfrak{E}_{\mathcal{A}}$ which goes on \mathcal{A} 's side of the black hole (i.e. passes through $\varphi = 0$), or a pair of disconnected extremal surfaces, $\mathfrak{E}_{\mathcal{A}^c}$ and the bifurcation surface \mathcal{H} of the event horizon $r = r_h$.

Now let us consider what happens when the causal wedge has a hole, as in figure 3. A connected spacelike²⁹ surface anchored at $\partial\mathcal{A}$ would then either have to stay at larger radial position than the connected part of $\Xi_{\mathcal{A}}$ (which is ruled out by the homology requirement), or it must pass around the black hole, in which case it must enter the causal wedge. To justify the latter more formally, in order to pass through $\varphi = 0$ and remain outside the causal wedge there, $\mathfrak{E}_{\mathcal{A}}$ would have to attain smaller radial value than the deepest reach of the disconnected part of $\Xi_{\mathcal{A}}$, so it would have to pass through the radial region traversed by the causal wedge caustics (\mathcal{C}^{\pm} in figure 3). Since this radial region is contained within the causal wedge at all angles, this suffices to guarantee the passage though the causal wedge. But since that would violate Property 8, no such connected extremal surface $\mathfrak{E}_{\mathcal{A}}$ homologous to \mathcal{A} can exist. \square

Comparing this result to the behaviour of spacelike geodesics in BTZ (cf. the right panel of figure 6), the non-existence of requisite extremal surfaces in higher dimensions might seem rather surprising — in fact, we are not aware of this effect having been noticed previously. Given that these surfaces exist only for sufficiently small $\varphi_{\mathcal{A}}$, and reach deeper into the bulk as $\varphi_{\mathcal{A}}$ increases, one might naturally wonder what happens as we try to push this deepest reach $r_{\mathfrak{E}}$ further towards the horizon. This is examined in detail in [26], with the curious result that the corresponding $\varphi_{\mathcal{A}}$ does not behave monotonically; instead it oscillates between certain maximal and minimal values. Even more curiously, $\varphi_{\mathcal{A}}(r_{\mathfrak{E}})$ oscillates infinitely many times as $r_{\mathfrak{E}} \rightarrow r_h$, exhibiting a self-similar behaviour. In other

²⁹ Since the bulk geometry is static, in this case of an eternal black hole we can in fact w.l.o.g. take the extremal surface (which is anchored at constant t on the boundary) to lie at constant t in the bulk. (In more complicated geometries such as discussed in [35] where there could be more general extremal surfaces, the homology requirement would no longer rule out $\mathfrak{E}_{\mathcal{A}^c}$ [26], so the present consideration would be irrelevant.)

words, for a fixed region \mathcal{A} with $\varphi_{\mathcal{A}}$ in a certain range, there are in fact infinitely many extremal surfaces (all outside the causal wedge) anchored on $\partial\mathcal{A}$ and homologous to \mathcal{A} . But conversely for large enough $\varphi_{\mathcal{A}}$, there are none.

Let us now ask what are the implications of these results for the entanglement entropy. As discussed in [26], we learn that the entanglement entropy cannot be a smooth function of $\varphi_{\mathcal{A}}$. At some critical $\varphi_{\mathcal{A}}$, the relevant surface which determines entanglement entropy switches from the $\mathfrak{E}_{\mathcal{A}}$ to the $\mathfrak{E}_{\mathcal{A}^c} + \mathcal{H}$ family, at which point $S(\varphi_{\mathcal{A}})$ has a kink and saturates onto a plateau. This phenomenon has been described as holographic entanglement plateau in [26], whose figure 9 shows the various critical curves for $\varphi_{\mathcal{A}}(r_h)$ in the case of Schwarzschild-AdS₅ black holes. We note that these authors choose to display the result in terms of $\alpha = \frac{1}{\pi}(\varphi_{\mathcal{A}} - \sin\varphi_{\mathcal{A}} \cos\varphi_{\mathcal{A}})$; for the casual wedges this is a rescaled version of our figure 5. We must emphasize that the critical point $\varphi_{\mathcal{A}}(r_h)$ for the entanglement plateau transition of [26] is only weakly determined by the casual wedge topology (indeed the presence of such an effect has been widely argued for since the inception of the holographic entanglement entropy proposal [21, 36]). This has to do with the fact that while non-trivial topology in the causal wedge is sufficient for the entanglement plateaux to develop, it certainly is not necessary. To ascertain the onset of the plateau phenomenon one requires detailed dynamics of minimal surfaces, and indeed as in any first order phase transition the exchange of dominance between the $\mathfrak{E}_{\mathcal{A}}$ to the $\mathfrak{E}_{\mathcal{A}^c} + \mathcal{H}$ families occurs before it is mandated by causal wedge topology.

5 Discussion

We have examined global properties of causal wedges and related constructs, proving a number of useful relational statements valid in arbitrary causal asymptotically AdS bulk spacetime. While from the CFT standpoint, the causal wedge might have been hitherto viewed as a rather esoteric construct, the preceding section demonstrated that its properties bear on more familiar CFT quantities such as the entanglement entropy. Nevertheless, the ultimate goal of this exercise is to use these properties to try to propose or actually construct the CFT dual of the causal wedge or associated quantities. The relational statements we discuss in section 3 are physically quite reasonable, and indeed unsurprising; the most curious feature, analyzed in section 2, is the topologically non-trivial nature of the causal wedge $\diamond_{\mathcal{A}}$ for simple regions \mathcal{A} .

Bolstered by the remarkable ease with which we have been able to produce ‘holes in the causal wedge’, the reader might well wonder whether this is perhaps the generic situation, namely whether in *any* bulk geometry satisfying the genericity condition, we could identify *some* region \mathcal{A} for which $\Xi_{\mathcal{A}}$ would have disconnected components. This would have strong implications for any putative CFT dual.

Since for the Schwarzschild-AdS _{$d+1$} family of solutions, the case which was guaranteed to have disconnected $\Xi_{\mathcal{A}}$ was the one with maximal $\mathcal{A} = \mathbf{S}^{d-1} \setminus i^0$, i.e. $\varphi_{\mathcal{A}} = \pi^-$, let us examine this situation for general bulk geometries. What is the generic form of $\Xi_{\mathcal{A}}$ and $\diamond_{\mathcal{A}}$? Note that domain of dependence³⁰ $\diamond_{\mathcal{A}}$ in this case coincides with the AdS Poincaré

³⁰ Typically domain of dependence \diamond_{Σ} is defined for a closed achronal surface Σ ; whereas in the present

wedge boundary, and in particular includes the Poincaré wedge spatial infinity vertex i^0 . Moreover the tips $q^{\wedge, \vee}$ of $\diamond_{\mathcal{A}}$ correspond to i^{\pm} of the Poincaré wedge. This in turn implies that $\diamond_{\mathcal{A}}$ does include i^0 , and since this point lies both on $\partial_-(\diamond_{\mathcal{A}})$ and $\partial_+(\diamond_{\mathcal{A}})$ (joined to $q^{\wedge, \vee}$ by the boundary null geodesics), it is contained in $\Xi_{\mathcal{A}}$. Now the question is, is $\Xi_{\mathcal{A}}$ just that one point i^0 , or does it reach into the bulk? In other words, are there points on bulk Cauchy slice anchored at $t_{\mathcal{A}}$ which don't lie in $\diamond_{\mathcal{A}}$?

For black holes the answer is clearly yes; but even for causally trivial spacetimes, one might naively expect that the answer is generically *yes*, with pure AdS being the only exception. The motivation for this expectation is as follows: consider a near-boundary point p just radially in from i^0 . In AdS, $J^{\pm}(p)$ intersects the boundary within $\diamond_{\mathcal{A}}$ only near $q^{\wedge, \vee}$. When we excite the CFT state by a bulk deformation (satisfying the null energy condition), these bulk geodesics are time-delayed [2, 3] so that they don't make it to $\diamond_{\mathcal{A}}$. Hence p would seem not to lie in $\diamond_{\mathcal{A}}$; but since $\diamond_{\mathcal{A}}$ does contain the entire boundary sphere, it should have a hole around p . This would lead us to expect that $\Xi_{\mathcal{A}}$ generically reaches into the bulk (presumably in some teardrop-like shape anchored on the single boundary point i^0).

However, there is a subtlety with this reasoning, having to do with the presence of caustics in $\partial\diamond_{\mathcal{A}}$. It is easy to see that radial null geodesics are time-delayed by some finite amount when traversing the bulk with some gravitational potential well. However, these are irrelevant, because generically they exit $\partial\diamond_{\mathcal{A}}$ and enter inside the causal wedge at a caustic point formed by null geodesics with ‘angular momentum’³¹ $\pm\varepsilon$, for some $\varepsilon \ll 1$. The geodesics which we need to focus on are the ones with maximal angular momentum. These are repelled by the centrifugal potential and as such stay close to the boundary feeling thereby less influenced by the deformation in the core (IR) region of the geometry.³² This suggests that the relevant geodesics to consider are the near-boundary geodesics with angular momentum $\ell = \pm(1 - \varepsilon)$.

Null geodesics through deformed AdS (for a static spherically symmetric metric corresponding to a “star” geometry) were examined by [10] in the context of bulk cone singularities. Let us consider a null congruence emanating from a specified point q^{\vee} on the boundary. In pure AdS, each null geodesic in this congruence would reach the boundary at the antipodal point, with³³ $\Delta\varphi = \pi$ and $\Delta t = \pi$. When the geometry has a gravitational potential well, the geodesics exhibit the usual time-delay and light-bending effect, which leads to a deformation of the future endpoint $(\Delta\varphi, \Delta t)$ of each geodesic, parameterized

case \mathcal{A} is open. However, since \diamond_{Σ} is defined as the set of all points from which every *timelike* curve intersects Σ , it should still be true in the present case that \diamond_{Σ} is closed.

³¹ Angular momentum along geodesics is a well-defined constant of motion only in presence of orbital Killing field; but as a warm-up to address part of the genericity question, we will focus on spherically symmetric spacetimes in this discussion.

³² We saw this behaviour in figure 3, but one can also verify this explicitly in more general static spherically symmetric examples: the larger angular momentum geodesics are less time delayed than the smaller angular momentum ones (including the radial one), so they will remain on $\partial\diamond_{\mathcal{A}}$ longer, covering up the smaller-angular momentum geodesics under a seam of caustics.

³³ Note that [10] considered equatorial geodesics with polar angle $\varphi \in [0, 2\pi)$, whereas here we are using $\varphi \in [0, \pi]$ to represent the azimuthal angle. However, as explained in section 2.1 due to spherical symmetry the geodesic equations are of course identical, so we can conflate the two for direct comparison.

by the reduced angular momentum $\ell \in (0, 1)$. A parametric plot of $\Delta t(\Delta\varphi)$ of geodesic endpoints for typical static spherically symmetric deformations of AdS were plotted in [10] (cf., their figure 5); the slope of the curve (parameterized by ℓ) in the $(\Delta\varphi, \Delta t)$ plane representing null geodesic endpoints is simply ℓ . This means that in general, near-boundary null geodesics reach the boundary at a more spatially than temporally shifted endpoint, $\Delta\varphi > \Delta t$. This would seem to imply that the $\pm\ell$ geodesics in the same (future or past) congruence intersect each other ‘before’ (i.e. closer to $\diamond_{\mathcal{A}}$) than they intersect the corresponding ℓ geodesic from the other congruence.

One can check these expectations explicitly in a given deformed bulk geometry. For example, consider the bulk metric (2.2) with $g(\rho) = 1 - \nu \cos^4 \rho$ with $\nu \ll 1$. This is a causally trivial asymptotically AdS geometry with a gravitational potential well given by ν , with pure AdS corresponding to $\nu = 0$. The causal wedge indeed remains topologically trivial, with a seam of caustics reaching all the way to i^0 . In fact, this is to be expected from the black hole result of section 2.1: the shape of (the connected component of) $\Xi_{\mathcal{A}}$ for large enough $\varphi_{\mathcal{A}}$ should be determined only by the asymptotic geometry, and should therefore be the same for corresponding to either a black hole or a star of the same mass. From figure 4 we see that as $\varphi_{\mathcal{A}} \rightarrow \pi$ (cf. the purple curves on the left), $\Xi_{\mathcal{A}}$ does *not* approach a tear-drop shape. In particular, without the disconnected component around the horizon, $\Xi_{\mathcal{A}}$ would indeed simply retract to i^0 as $\varphi_{\mathcal{A}} \rightarrow \pi$. So we expect that in order to obtain causal wedges with non-trivial topology, the bulk spacetime has to be sufficiently deformed from pure AdS.

Said another way, there is an interesting tension between two competing effects. Had the time-delay effect been the only operative feature, we would have concluded that holes in the causal wedge would have been generic. On the other hand, if centrifugal repulsion which forces attention on near-boundary geodesics were to be the primary effect, there would be no room for non-trivial topology. It is then clear that the natural place to look for non-trivial causal wedge topology is when the two effects are competing, and indeed our prototypical example of Schwarzschild-AdS provides one such setting. Naively then, we can abstract from the above discussion an essential requirement for the causal wedge to develop holes: *the spacetime must admit null circular orbits.*³⁴ For it is in this case that there is some non-trivial interplay between the two effects discussed above and they precisely offset each other at the circular orbit. Indeed we saw that the critical transition point for Schwarzschild-AdS was effectively determined by the null circular orbit in section 2.1. The justification provided there leads us to conjecture that in spherically symmetric spacetimes the presence of the null circular orbits is a necessary condition for non-trivial causal wedge topology. It would be useful to prove this statement rigorously; we hope to return to this interesting problem in the near future.

One lesson of our explorations which we wish to emphasize is that while studies of global AdS contain the Poincaré AdS case (as a limit), the converse is not true. Not only is the global case much richer (as can be expected already from the metric being more

³⁴ Of course, this can only be a requirement for spherically symmetric spacetimes, where such orbits may exist in the first place. In absence of any symmetries, we would have to formulate a more robust criterion.

complicated), but it gives us novel insight into important observables. Over the years we have learnt many interesting lessons by examining field theories on compact spatial geometries and we believe that there is indeed much more to be learned from similar explorations.

Acknowledgments

It is a pleasure to thank Matt Headrick and Harvey Reall useful discussions. MR and VH would like to thank CERN, ITF Amsterdam, ICTP and Crete Center for Theoretical Physics for hospitality during the course of this project. MR would like to thank the organizers of the “Relativistic fluid dynamics and the gauge gravity duality” workshop held at Technion, Haifa and those of the “2nd Mediterranean Conference on Classical And Quantum Gravity”, Veli Losinj, Croatia for their hospitality during the concluding stages of this project. ET would like to thank the NBI Copenhagen and Physics dept. of the University of Turin for hospitality during part of this work. VH and MR are supported in part by the the STFC Consolidated Grant ST/J000426/1.

References

- [1] G.T. Horowitz and N. Itzhaki, *Black holes, shock waves and causality in the AdS/CFT correspondence*, *JHEP* **02** (1999) 010 [[hep-th/9901012](#)] [[INSPIRE](#)].
- [2] E. Woolgar, *The Positivity of energy for asymptotically anti-de Sitter space-times*, *Class. Quant. Grav.* **11** (1994) 1881 [[gr-qc/9404019](#)] [[INSPIRE](#)].
- [3] S. Gao and R.M. Wald, *Theorems on gravitational time delay and related issues*, *Class. Quant. Grav.* **17** (2000) 4999 [[gr-qc/0007021](#)] [[INSPIRE](#)].
- [4] D.N. Kabat and G. Lifschytz, *Gauge theory origins of supergravity causal structure*, *JHEP* **05** (1999) 005 [[hep-th/9902073](#)] [[INSPIRE](#)].
- [5] V. Balasubramanian and S.F. Ross, *Holographic particle detection*, *Phys. Rev. D* **61** (2000) 044007 [[hep-th/9906226](#)] [[INSPIRE](#)].
- [6] J. Louko, D. Marolf and S.F. Ross, *On geodesic propagators and black hole holography*, *Phys. Rev. D* **62** (2000) 044041 [[hep-th/0002111](#)] [[INSPIRE](#)].
- [7] J.P. Gregory and S.F. Ross, *Looking for event horizons using UV/IR relations*, *Phys. Rev. D* **63** (2001) 104023 [[hep-th/0012135](#)] [[INSPIRE](#)].
- [8] D. Marolf, *States and boundary terms: Subtleties of Lorentzian AdS/CFT*, *JHEP* **05** (2005) 042 [[hep-th/0412032](#)] [[INSPIRE](#)].
- [9] V.E. Hubeny, M. Rangamani and S.F. Ross, *Causal structures and holography*, *JHEP* **07** (2005) 037 [[hep-th/0504034](#)] [[INSPIRE](#)].
- [10] V.E. Hubeny, H. Liu and M. Rangamani, *Bulk-cone singularities & signatures of horizon formation in AdS/CFT*, *JHEP* **01** (2007) 009 [[hep-th/0610041](#)] [[INSPIRE](#)].
- [11] R. Bousso, S. Leichenauer and V. Rosenhaus, *Light-sheets and AdS/CFT*, *Phys. Rev. D* **86** (2012) 046009 [[arXiv:1203.6619](#)] [[INSPIRE](#)].
- [12] 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](#)] [[INSPIRE](#)].

- [13] V.E. Hubeny and M. Rangamani, *Causal Holographic Information*, *JHEP* **06** (2012) 114 [[arXiv:1204.1698](#)] [[INSPIRE](#)].
- [14] R. Bousso, B. Freivogel, S. Leichenauer, V. Rosenhaus and C. Zukowski, *Null Geodesics, Local CFT Operators and AdS/CFT for Subregions*, [arXiv:1209.4641](#) [[INSPIRE](#)].
- [15] S. Ryu and T. Takayanagi, *Aspects of Holographic Entanglement Entropy*, *JHEP* **08** (2006) 045 [[hep-th/0605073](#)] [[INSPIRE](#)].
- [16] V.E. Hubeny, M. Rangamani and T. Takayanagi, *A Covariant holographic entanglement entropy proposal*, *JHEP* **07** (2007) 062 [[arXiv:0705.0016](#)] [[INSPIRE](#)].
- [17] V.E. Hubeny, *Extremal surfaces as bulk probes in AdS/CFT*, *JHEP* **07** (2012) 093 [[arXiv:1203.1044](#)] [[INSPIRE](#)].
- [18] P.L. Ribeiro, *Algebraic holography in asymptotically simple, asymptotically AdS spacetimes*, *Prog. Math.* **251** (2007) 253 [[hep-th/0502096](#)] [[INSPIRE](#)].
- [19] P.L. Ribeiro, *Structural and dynamical aspects of the AdS/CFT correspondence: A Rigorous approach*, [arXiv:0712.0401](#) [[INSPIRE](#)].
- [20] V.E. Hubeny, M. Rangamani and E. Tonni, *Thermalization of Causal Holographic Information*, *JHEP* **05** (2013) 136 [[arXiv:1302.0853](#)] [[INSPIRE](#)].
- [21] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Phys. Rev. Lett.* **96** (2006) 181602 [[hep-th/0603001](#)] [[INSPIRE](#)].
- [22] B. Freivogel and B. Mosk, *Properties of Causal Holographic Information*, [arXiv:1304.7229](#) [[INSPIRE](#)].
- [23] G. Galloway, K. Schleich, D. Witt and E. Woolgar, *Topological censorship and higher genus black holes*, *Phys. Rev. D* **60** (1999) 104039 [[gr-qc/9902061](#)] [[INSPIRE](#)].
- [24] J.J. Friess, S.S. Gubser, G. Michalogiorgakis and S.S. Pufu, *Expanding plasmas and quasinormal modes of anti-de Sitter black holes*, *JHEP* **04** (2007) 080 [[hep-th/0611005](#)] [[INSPIRE](#)].
- [25] H. Araki and E. Lieb, *Entropy inequalities*, *Commun.Math.Phys.* **18** (1970) 160.
- [26] V.E. Hubeny, H. Maxfield, M. Rangamani and E. Tonni, *Holographic entanglement plateaux*, *JHEP* **08** (2013) 092 [[arXiv:1306.4004](#)] [[INSPIRE](#)].
- [27] A.D. Helfer, *Black holes reconsidered*, [arXiv:1105.1980](#) [[INSPIRE](#)].
- [28] P.T. Chrusciel and G.J. Galloway, *‘Nowhere’ differentiable horizons*, *Commun. Math. Phys.* **193** (1998) 449 [[gr-qc/9611032](#)] [[INSPIRE](#)].
- [29] J.K. Beem and A. Krolak, *Cauchy horizon endpoints and differentiability*, [gr-qc/9709046](#) [[INSPIRE](#)].
- [30] B. Freivogel, J. McGreevy and S.J. Suh, *Exactly Stable Collective Oscillations in Conformal Field Theory*, *Phys. Rev. D* **85** (2012) 105002 [[arXiv:1109.6013](#)] [[INSPIRE](#)].
- [31] P. Figueras, V.E. Hubeny, M. Rangamani and S.F. Ross, *Dynamical black holes and expanding plasmas*, *JHEP* **04** (2009) 137 [[arXiv:0902.4696](#)] [[INSPIRE](#)].
- [32] J. Abajo-Arastia, J. Aparicio and E. Lopez, *Holographic Evolution of Entanglement Entropy*, *JHEP* **11** (2010) 149 [[arXiv:1006.4090](#)] [[INSPIRE](#)].
- [33] V.E. Hubeny and H. Maxfield, *Holographic probes of collapsing black holes*, work in progress.

- [34] A.C. Wall, *Maximin Surfaces and the Strong Subadditivity of the Covariant Holographic Entanglement Entropy*, [arXiv:1211.3494](#) [[INSPIRE](#)].
- [35] B. Freivogel, V.E. Hubeny, A. Maloney, R.C. Myers, M. Rangamani et al., *Inflation in AdS/CFT*, *JHEP* **03** (2006) 007 [[hep-th/0510046](#)] [[INSPIRE](#)].
- [36] M. Headrick and T. Takayanagi, *A Holographic proof of the strong subadditivity of entanglement entropy*, *Phys. Rev. D* **76** (2007) 106013 [[arXiv:0704.3719](#)] [[INSPIRE](#)].