

Light-sheets and AdS/CFT

 Raphael Bousso,^{1,2} Stefan Leichenauer,³ and Vladimir Rosenhaus^{1,2}
¹*Center for Theoretical Physics and Department of Physics, University of California, Berkeley, California 94720, USA*
²*Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*
³*California Institute of Technology, Pasadena, California 91125, USA*

(Received 11 June 2012; published 24 August 2012)

One may ask whether the conformal field theory (CFT) restricted to a subset b of the anti-de Sitter (AdS) boundary has a well-defined dual restricted to a subset $H(b)$ of the bulk geometry. The Poincaré patch is an example, but more general choices of b can be considered. We propose a geometric construction of H . We argue that H should contain the set C of causal curves with both endpoints on b . Yet H should not reach so far from the boundary that the CFT has insufficient degrees of freedom to describe it. This can be guaranteed by constructing a superset L of H from light-sheets off boundary slices and invoking the covariant entropy bound in the bulk. The simplest covariant choice is $L = L^+ \cap L^-$, where L^+ (L^-) is the union of all future-directed (past-directed) light-sheets. We prove that $C = L$, so the holographic domain is completely determined by our assumptions: $H = C = L$. In situations where local bulk operators can be constructed on b , H is closely related to the set of bulk points where this construction remains unambiguous under modifications of the CFT Hamiltonian outside of b . Our construction leads to a covariant geometric renormalization-group flow. We comment on the description of black hole interiors and cosmological regions via AdS/CFT.

 DOI: [10.1103/PhysRevD.86.046009](https://doi.org/10.1103/PhysRevD.86.046009)

PACS numbers: 11.25.Tq

I. INTRODUCTION

In the study of the Lorentzian AdS/CFT correspondence [1–6], one may consider the boundary theory defined on a proper subset of the global conformal boundary. In this case, it is natural to expect that the bulk dual spacetime manifold may be extendable. That is, the theory on the subset should describe less than the maximally extended global bulk solution.

The most common example is the conformal field theories (CFT) on Minkowski space, $\mathbf{R}^{d-1,1}$, which is dual to the Poincaré patch of the bulk, with metric

$$ds^2 = R^2 \frac{-dt^2 + d\vec{x}_{d-1}^2 + dz^2}{z^2}. \quad (1.1)$$

These bulk and boundary regions are shown as subsets of global anti-de Sitter (AdS) in Fig. 1(a). The Poincaré patch seems a “natural” choice for the bulk, at least in the above warped-product coordinate system, where the boundary corresponds to $z = 0$ and the Poincaré patch to $z > 0$. However, there is nothing special about this choice of coordinates. One could easily write down coordinates that cover a larger or smaller portion of the bulk that is bounded by the same portion of the conformal boundary as $z \rightarrow 0$ [see Fig. 1(b)]. So what selects the Poincaré patch as the bulk dual of the CFT on Minkowski space?

The bulk dual region should be well defined not only in the vacuum, but for arbitrary states in the

semiclassical regime, perturbatively in $1/N$. Deep in the bulk, the metric will not be that of empty AdS space. In general, all Killing symmetries will be broken, so they cannot be used to pick out a preferred bulk region. For example, consider a global bulk solution corresponding to a pair of neutron stars orbiting around the origin of the standard global coordinates. At what time (say, along its worldline) does each star enter and exit the bulk region dual to the CFT on $\mathbf{R}^{d-1,1}$?




FIG. 1. (a) The Poincaré patch of AdS, with the usual time slicing in the coordinates of Eq. (1.1). (b) Time slices of an arbitrary bulk coordinate system that covers the same near-boundary region as the Poincaré patch but a different region far from the boundary. This illustrates that there is no preferred coordinate system that would uniquely pick out a region described by the boundary, particularly if the bulk is not in the vacuum state.

Another well-known example is the maximally extended Schwarzschild-AdS black hole. The global conformal boundary consists of two disconnected copies of $\mathbf{R} \times \mathbf{S}^{d-1}$. In the Hartle-Hawking state (the Euclidean vacuum), the two components can be thought of as slices in a single complex manifold and are related by analyticity. However, one can consider more general states, for example by adding neutron stars near the left boundary (far from the black hole), and white dwarfs near the right boundary. Restricting attention to the CFT living on the left boundary, one would expect it to encode the nearby neutron stars, but not the white dwarfs on the far side of the black hole. But what is the basis of this expectation? And where does the reach of the left theory end: at the black hole horizon, or somewhere inside the black hole/white hole regions? Again, there should be an answer to this question that does not rely on special bulk symmetries or coordinate choices.

There are many other possible choices of subsets b of the global boundary, some of which are shown in Fig. 2.¹ Let us suppose that the CFT on b describes some portion $H(b)$ of the bulk. We will call $H(b)$ the holographic domain of b . It should be possible to construct this bulk region geometrically from b . The goal of this paper is to provide such a construction.²

Our strategy will be to bound the bulk dual region $H(b)$ by a subset $C(b)$ and a superset $L(b)$:

$$C(b) \subset H(b) \subset L(b). \quad (1.2)$$

¹We consider only subsets b which, viewed as manifolds in their own right, are globally hyperbolic and have the same number of spacetime dimensions as the global boundary. Otherwise the initial value problem of the CFT would be ill defined. If b is a proper subset of the global boundary, then it is not obvious that the CFT on b must have a bulk dual, and we do not set out to prove this or establish under which conditions a bulk dual exists. The question we seek to address is: if b did have a bulk dual region, what would it be? In all cases, a semiclassical bulk dual can be assigned only to some subset of CFT states (excluding, for example, the equal superposition of two global CFT states corresponding to different bulk metrics), and only perturbatively in $1/N$. In the case of a small diamond, the bulk dual is the AdS-Rindler patch, so one would expect the geometric states to have a thermal character.

²In the special case that b is a diamond, motivated by the Ryu-Takayanagi proposal [7–9] for computing holographic entanglement entropy, one could try constructing $H(b)$ out of minimal surface which start and end on b . We will not explore this approach here, however it would be interesting to study how it relates to our construction. We will note however that Ryu-Takayanagi only applies for static situations; for time-dependent situations one must use the proposal [10] of (minimal) extremal surfaces, which has been less well established. Additionally, as we will find in Sec. VI, allowing sufficient modifications of the boundary theory makes the bulk ambiguous in large regions. In this sense, the bulk dual naturally constructed from extremal surfaces would generally be too large.

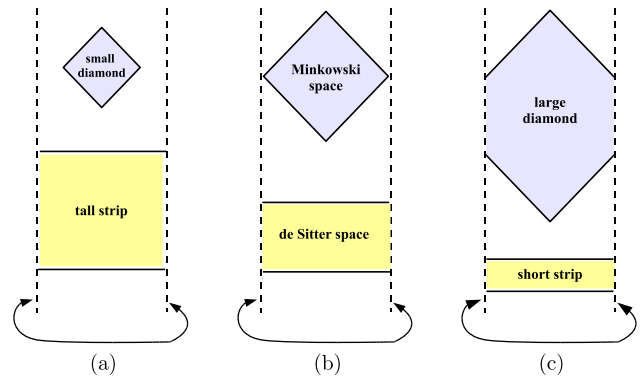


FIG. 2 (color online). The boundary of AdS; the dashed lines should be identified. Examples of globally hyperbolic subsets b are shown shaded. A causal diamond is a set of the form $I^-(q) \cup I^+(p)$, where q is boundary event in the future of the boundary event p . Let τ be the time along a geodesic from p to q in the Einstein static universe of unit radius ($ds^2 = -dt^2 + d\Omega_{d-1}^2$). With $\tau = 2\pi$, the causal diamond is the boundary of the Poincaré patch. A causal diamond with $\tau < 2\pi$ ($\tau > 2\pi$) is called “small” (“large”). An open interval (t_1, t_2) with $t_2 - t_1 < \pi$ ($t_2 - t_1 > \pi$) is called “short strip” (“tall strip”).

We choose $C(b)$ to be the minimum region the CFT needs to describe to be consistent with bulk causality and locality properties near the boundary. The set $L(b)$ will be constructed from light-sheets in a way that ensures that the bulk dual does not have more degrees of freedom (higher maximum entropy) than the CFT. We will then show that under some additional assumptions on $L(b)$, all three sets are equal. This fully determines $H(b)$ subject to the stated assumptions.

In Sec. II, we propose a lower bound on the holographic domain. We argue that $H(b) \supset C(b)$, where $C(b) = J^+(b) \cap J^-(b)$ is the set of bulk points which lie both in the causal past and in the causal future of the conformal boundary region b . Otherwise, the boundary would fail to describe regions that can be explored by a bulk observer localized near the boundary.

In Sec. III, we propose an upper bound on the holographic domain: $H(b) \subset L(b)$, determined by the requirement that the boundary theory should contain enough degrees of freedom to describe the bulk. The number of CFT degrees of freedom is controlled by the area A of the regulated boundary [11]. In order to ensure that the CFT can describe the holographic domain H , H must be contained in a bulk region that satisfies an entropy bound $S \lesssim A$. In Sec. III A, we review the covariant entropy bound, which states that the relevant entropy lies on *light-sheets* emanating from the boundary area. There are no generally valid bounds on the entropy at equal time, so light-sheets should play a preferred role in any holographic duality. A light-sheet is a null hypersurface generated by nonexpanding light rays orthogonal to a codimension-2 surface.

The boundary is of codimension-1, so the construction of a bulk region out of light-sheets off the boundary first

requires a foliation of the boundary b into time slices (which are codimension-2 surfaces from the bulk viewpoint). In Sec. III B, we examine whether it is nevertheless possible to define a covariant, slicing-independent upper bound $L(b)$ from light-sheets.³

Given a boundary time slicing, one may consider the union $L^+(b)$ of all future-directed light-sheets (one from each slice), or the union $L^-(b)$ of all past-directed light-sheets. However, both of these sets depend on the slicing of b , and the same is true for the union $L^+(b) \cup L^-(b)$. Thus, neither defines an upper bound on $H(b)$ covariantly.

We then consider the intersection $L(b) \equiv L^+(b) \cap L^-(b)$ as a candidate for an upper bound on $H(b)$. There are two apparent problems with this choice. The first is that $L(b)$, too, would appear to depend on the time slicing of the boundary set b . Second, it is easy to see that $L(b) \subset C(b)$. This would conflict with our expectation that $C(b) \subset H(b) \subset L(b)$, unless it can be shown that $L(b) = C(b)$. That is, the consistency of our arguments requires that any point that lies on a causal curve that begins and ends on the boundary region b must also lie on both a past- and a future-directed light-sheet emanating from b , independently of the time slicing of b .

We prove this nontrivial result in Sec. IV. As a corollary, the slicing independence of $C(b)$ implies the slicing independence of $L(b)$. Assuming this is the correct choice of $L(b)$, it follows that the holographic domain $H(b)$ is completely determined:

$$H(b) = C(b) = L(b). \quad (1.3)$$

The description of $H(b)$ in terms of light-sheets allows us to define a holographic bulk renormalization-group (RG) flow, and the equivalence with $C(b)$ makes the RG flow manifestly covariant. Combining the focusing theorem with the covariant entropy bound guarantees that the number of degrees of freedom is strictly nonincreasing along the flow. We discuss this construction in Sec. V.

In Sec. VI we consider natural definitions of the bulk dual region from the perspective of the field theory, which contains nonlocal operators that can probe deeply into the bulk. This approach is less general than our geometric approach, because the construction of such operators is not known for arbitrary b . In cases where it is known we find that a region \bar{H} , very similar to $H(b)$, is picked out as the region with an unambiguous bulk interpretation.

Section VII is largely independent of the rest of this paper. We consider the extent to which AdS/CFT adds to

³Light-sheets were used in Ref. [12] to determine the holographic domain of (effectively) half of the global boundary. This region is not globally hyperbolic and the theorem of Sec. IV does not apply; but the division of the boundary selects for a preferred slicing. Recently, Hubeny has examined which bulk regions are explored by extremal surfaces of various dimensions in static situations; she finds, as we do for the holographic domain H , that black hole interiors are never included [13].

our understanding of quantum gravity in regions with dominant self-gravity, such as black hole interiors and cosmological regions. We describe a thought-experiment involving the formation and evaporation of (smaller) black holes in such regions, and we argue that its description requires a nonperturbative bulk theory.

II. THE CAUSALLY CONNECTED REGION C

In this section, we argue that the CFT on the boundary portion b must at least describe the set $C(b)$ of bulk points that are both in the past and in the future of b : $H(b) \supset C(b)$.

Consider a bulk excitation very close to the boundary, e.g., at $\frac{\pi}{2} - \rho = \epsilon \ll 1$ in global coordinates,

$$ds^2 = \frac{R^2}{\cos^2 \rho} (-d\tau^2 + d\rho^2 + \sin^2 \rho d\Omega_{d-1}^2). \quad (2.1)$$

Such an excitation is represented on the boundary by excitations with support on a region of size ϵ , at the same transverse position [11]. Here we take the boundary theory to live on a unit sphere, but this property is essentially local. It holds in Poincaré as well as global coordinates. One expects, therefore, that it should hold for any boundary region b , as long as ϵ is much smaller than the characteristic temporal and spatial size of b . This implies that the boundary theory on b must describe at least the state and the dynamics of a bulk region sufficiently close to b . We will exploit the fact that this near-boundary region, in turn, is causally connected to the larger bulk region $C(b)$ to show that the CFT must at least describe $C(b)$.

Consider now a family of bulk observers localized at $\frac{\pi}{2} - \rho = \epsilon \ll 1$. They will require a finite proper acceleration of order R^{-1} , where R is the AdS curvature radius. The proper time for which this acceleration must be maintained is

$$\Delta t_{\text{proper}} = \frac{R}{\epsilon} \Delta \tau. \quad (2.2)$$

This diverges in the limit as $\epsilon \rightarrow 0$, but we will only need to consider the case of small but finite ϵ . Thus there appears to be no obstruction, in principle, to the existence of such bulk observers. What region can they explore? We will treat the observers as a collective, imagining that they densely fill Cauchy surfaces of the boundary b , moved into the bulk by ϵ . Such observers can receive signals from the causal past of b in the bulk, and they can send signals to the causal future of b . However, they cannot determine the state in the entirety of either of these regions without making additional assumptions.

The only region that can be actively explored and manipulated by observers near the boundary is the intersection of the past and the future of b ,

$$C(b) \equiv J^+(b) \cap J^-(b). \quad (2.3)$$

This region can be probed by preparing local probes at an early time, allowing them to travel deeper into the bulk and then back to the observer at a late time. The outcome of such an experiment is completely determined by the state in $C(b)$, by causality. And conversely, local fields at any point in $C(b)$ can be manipulated by such an experiment. Since the experiments of these near-boundary observers are described by the CFT, then for a consistent duality to hold the CFT must describe at least the region $C(b)$:

$$C(b) \subset H(b). \quad (2.4)$$

If the near boundary region can both probe and be affected by $C(b)$, then the same must be true for the boundary theory itself.

By contrast, knowledge of the state in $J^-(b) - J^+(b)$, say, is not necessary in order to compute the dynamics near the boundary or in any region that can be explored from the boundary. It is sufficient to specify initial conditions on the past boundary of $C(b)$ in the bulk. Conversely, since its past boundary need not be a Cauchy surface for $J^-(b) - J^+(b)$, the state in $J^-(b) - J^+(b)$ is not determined by the state in the region $C(b)$ which can be explored from the boundary.

III. THE LIGHT-SHEET REGION L

In this section we construct a region $L(b)$ which is guaranteed to be encodable in the CFT on the boundary portion b , in the sense that holographic entropy bounds guarantee that the maximum entropy of matter in $L(b)$ does not exceed the number of degrees of freedom of the CFT on b . For an exact duality to hold, $H(b)$ must be contained in some region with this property, and we will conjecture that $H(b) \subset L(b)$.

The holographic principle [14–17] is a universal relation between area and quantum information. This relation manifests itself empirically as the Covariant Entropy Bound [16]: Let A be the area of an arbitrary (open or closed) codimension-2 surface B , and let S be the entropy of the matter on any light-sheet ℓ of B :

$$S[\ell(B)] \leq \frac{A}{4}. \quad (3.1)$$

At the core of the covariant entropy bound is the notion of a light-sheet (Fig. 3). Light-sheets are null hypersurfaces generated by nonexpanding light rays orthogonal to the surface B . There are four orthogonal directions at every point, since the surface has two sides and we can consider past- and future-directed light rays. If the null energy condition is satisfied, at least two of these directions will have nonpositive expansion and thus give rise to light-sheets. The covariant entropy bound holds separately on each light-sheet. (For a review, see Ref. [18].)

A. Spacelike holography vs light-sheets

The holographic principle is *not* the statement that the entropy in any spatial region V is limited by the area of the




FIG. 3. The four null hypersurfaces orthogonal to a spherical surface B in Minkowski space. The two cones F_1, F_3 have negative expansion and hence correspond to light-sheets. The covariant entropy bound states that the entropy of the matter on each light-sheet will not exceed the area of B . The other two families of light rays, F_2 and F_4 , generate the skirts drawn in thin outline. Their cross-sectional area is increasing, so they are not light-sheets, and the entropy of matter on them is unrelated to the area of B .

surface bounding that region. This “spacelike” entropy bound follows from the covariant bound in certain special cases [16], but in general it is false. Counterexamples are easily found in cosmology, inside black holes, and even in weakly gravitating systems [18]. As a general statement valid in all spacetimes, one must not think of holography as effectively projecting out a spacelike direction. Holography projects along a null direction, just as it does in a conventional hologram.

To illustrate the falsehood of any spacelike notion of holography, consider the surface area of a collapsing star (Fig. 4). The area approaches zero near the singularity, but the entropy of the star starts out finite and cannot decrease. Thus, $S(V) \gg A$ at late times: a violation of the spacelike entropy bound. However, a light-sheet off of a late-time surface will not penetrate the whole star, so the covariant bound, Eq. (3.1), is upheld.

Another example is shown in Fig. 5: It is possible to surround any matter system with a surface of arbitrarily small area. (For weakly gravitating systems the surface will




FIG. 4 (color online). Penrose diagram of a collapsing star (shaded). At late times, the area of the star’s surface becomes very small (B). The enclosed entropy in the *spatial* region V stays finite, so that the spacelike entropy bound is violated. The covariant entropy bound avoids this difficulty because only future-directed light-sheets are allowed by the nonexpansion condition. L is truncated by the future singularity; it does not contain the entire star.




FIG. 5 (color online). (a) A square system in 2 + 1 dimensions, surrounded by a surface B of almost vanishing length A . The entropy in the enclosed *spatial* volume can exceed A . (b) Here the time dimension is projected out. The light-sheet of B intersects only with a negligible (shaded) fraction of the system, so the covariant entropy bound is satisfied.

consist of elements that are highly boosted with respect to the rest frame.) Thus, the spacelike entropy bound is violated. By contrast, the light-sheets off of such surfaces do not contain most of the system, evading violation of the covariant entropy bound. Details and further examples are given in Ref. [18].

Notice that the above counterexamples to a spacelike entropy bound can easily be embedded into an asymptotically AdS spacetime. In particular, consider a timelike hypersurface that is a direct product, $b = \mathbf{R} \times \mathbf{S}^2$. (For definiteness, we consider AdS_4 but our arguments hold in any number of dimensions.) In the limit where the spatial two-spheres are large, b encloses a very large spacetime region. We can think of b as a regulated version of the conformal boundary. The bulk spacetime it encloses is described by a conformal field theory living on b , with a UV cutoff on a length scale comparable to the AdS curvature radius [11].

Yet, the hypersurface b can also be foliated into two-dimensional slices which have arbitrarily small area; and each such slice bounds a global bulk slice. With this slicing, a naive “spacelike” interpretation of holography would seem to imply that the bulk can be described by a theory with an arbitrarily small Hilbert space. This interpretation is clearly incorrect, as the bulk can have arbitrarily large entropy in the limit where b approaches the boundary.

Another example is furnished by the eternal Schwarzschild-AdS black hole solution. Let us again consider the regulated global boundary given by the direct product of time with a sphere of fixed radius larger than the black hole, $b = \mathbf{R} \times \mathbf{S}^2$. As shown in Fig. 6, this sphere encloses an infinitely large spatial region that extends all the way to the second, disconnected conformal boundary on the far side of the black hole. This region contains arbitrarily large entropy: for example, a dilute gas of $n \rightarrow \infty$ photons can be added near the second boundary with negligible backreaction. Of course, our intuition tells




FIG. 6. An AdS-Schwarzschild black hole. A sphere on the regulated boundary encloses an infinitely large spatial region that extends all the way to the second, disconnected conformal boundary on the far side of the black hole.

us that the CFT that lives on the second boundary should be “responsible” for those photons. The only problem is that spacelike holography does *not* tell us this. In fact, it cannot possibly tell us any such thing. Black hole event horizons are global objects; for example, if we decide much later to add mass to the black hole, this will affect the location of the horizon at all earlier times. Thus, there is no local criterion that can prevent us from extending a spacelike holographic domain to the far side of the black hole.

As the above examples illustrate, the light-like nature of the holographic relation between entropy and area is crucial for making sense of AdS/CFT as a manifestation of the holographic principle. This relation is captured by the geometric construction of light-sheets.

B. Bounds on the holographic domain

As shown by Susskind and Witten [11], the CFT with a UV cutoff δ has a maximum entropy equal to the proper area of the spatial boundary of AdS in standard global coordinates. The latter is considered to be located in the bulk, a coordinate distance δ away from the true conformal boundary. This observation tells us that the CFT manifestly has the correct number of degrees of freedom demanded by the holographic principle, given by the area of the boundary, not by the size of the enclosed volume.

In this argument one assumes, of course, that the CFT describes no more and no less than the spacetime region “enclosed” in the regulated boundary. This is plausible in the case of global AdS with the standard slicing of the boundary into round spheres. The analysis of the previous subsection has shown, however, that this assumption is coordinate-dependent at best, and that it is ill-defined for cases such as the Poincaré patch, where the boundary slices do not “enclose” any particular region.

In order to claim more generally that the CFT uses an area’s worth of degrees of freedom to describe the bulk, we

must characterize some relevant bulk region, given the boundary. Note that the UV/IR connection is not the issue; by construction, the Susskind-Witten argument suffices to ensure that the maximum entropy on the CFT agrees with the area of the regulated boundary. This remains true for arbitrary foliations of the boundary, as long as the short-distance regulator in the CFT is imposed with respect to the chosen time slices. We can always remove the UV cutoff at the end and think of the CFT as living on the true conformal boundary.

The nontrivial question is how far from the boundary (how deep into the bulk) the CFT description is valid. If this region is taken to extend too far from the boundary, then the bulk entropy might be larger than the boundary area, and thus larger than the maximum entropy of the CFT. In that case, the CFT Hilbert space would be too small to capture the bulk physics. The only way to ensure that the bulk entropy is sufficiently small is to appeal to the covariant entropy bound. This is why the relevant bulk region must be related to the boundary by light-sheets. We will now explore possible concrete proposals for this relation.

Consider a timelike hypersurface b embedded in an asymptotically AdS and a foliation into spacelike hypersurfaces. We may view b as a spacetime in its own right, with one less spatial dimension than the AdS spacetime it is embedded in. In order for the theory living on b to be well defined, we require that b be globally hyperbolic and that each time slice be a Cauchy surface. Most relevant for the present discussion is the limiting case where b lies on the conformal boundary of global AdS: $b \subset I$. (b can be a proper subset of I , for example if b is the boundary of the Poincaré patch of AdS.) In this case, the metric of b is defined only up to conformal transformations. We will be interested only in situations where each time slice of b is “normal,” i.e., each slice t admits both a past-directed light-sheet ℓ_t^- and a future-directed light-sheet ℓ_t^+ . This is automatically the case for $b \subset I$, as we show in the Appendix.

Let us consider the regions $L^+(b)$ and $L^-(b)$. $L^+(b)$ is defined as the union over t of the future-directed light-sheets ℓ_t^+ of each slice. Similarly, $L^-(b)$ is the union of all past-directed light-sheets. Two natural looking possible bounds on the holographic domain are $H(b) \subset L^+(b)$ and $H(b) \subset L^-(b)$. At the fundamental level, there is no distinction between the past and the future, so $L^+(b)$ and $L^-(b)$ should play a symmetric role as bounds on the holographic domain. Let us therefore consider the candidate bounds $H(b) \subset L^+(b) \cup L^-(b)$, and $H(b) \subset L^+(b) \cap L^-(b)$. The former bound guarantees that every point in $L^+(b) \cup L^-(b)$ lies either on a future-directed or a past-directed light-sheet from some time slice of b . The latter bound is stronger: if $H(b) \subset L^+(b) \cap L^-(b)$, then every point in $H(b)$ lies on at least one past- and one future-directed light-sheet.

The choice of time slicing on b is a coordinate choice and so cannot have fundamental significance. Remarkably,

the set $L(b) \equiv L^+(b) \cap L^-(b)$ is indeed independent of how b is foliated, even though $L^+(b)$ and $L^-(b)$ individually do depend on the time slicing of b . In the following section we will prove a stronger theorem: $L(b) = C(b)$. (Recall that $C(b) \equiv J^+(b) \cap J^-(b)$ is the region causally connected to b , which we discussed in the previous section.) The fact that $L(b)$ is slicing independent follows as a corollary, since $C(b)$ is slicing independent by construction. The simplicity and slicing independence of $L(b)$ make it especially attractive, and we will see in Sec. V that it also leads to a useful formulation of holographic RG flow.

The other set that we were led to consider, the union $L^+(b) \cup L^-(b)$, is slicing dependent (see Fig. 7). This disqualifies it from further consideration, if we insist, as we do in this paper, that a unique upper bound on the holographic domain should be constructed from light-sheets off a single, arbitrary slicing of the boundary region b .

It is important to note that there are other, more complicated ways of constructing a bulk region covariantly from light-sheets. We will not analyze such sets here, but we mention some of them for completeness and future consideration. For example, to treat the possible time slicings in a democratic way, one could form the union, or the intersection, over all possible time slicings \mathcal{T} , of the sets $L^+(b) \cup L^-(b)$: $L'(b) = \bigcap_{\mathcal{T}} (L^+(b) \cup L^-(b))$ and $L''(b) = \bigcup_{\mathcal{T}} (L^+(b) \cup L^-(b))$. Clearly, $L(b) \subset L'(b) \subset L''(b)$. However, the restriction $H(b) \subset L''(b)$ seems too weak to ensure that the boundary theory has enough degrees of freedom to describe the bulk. For a given time slicing, it




FIG. 7 (color online). The union of all future-directed light-sheets, $L^+(b)$, coming off the usual slicing of the boundary Minkowski space (left) covers precisely the Poincaré patch (the wedge-shaped region that lies both in the future of the boundary point A and in the past of D). On the right, we show a different time slicing of the same boundary region. One of these slices is shown in blue in the bulk Penrose diagram (center); it curves up at B and down at C . The future-directed light-sheet coming off the portion of the slice near C is nearly the same as the future light cone of C (shown in red/short-dashed), which reaches far beyond the Poincaré patch to the far side of AdS. The bulk region covered by $L^+(b)$ will thus be nearly two Poincaré patches, consisting of the points that lie in the future of A but not in the future of D (long-dashed).

may not be the case that a given point in $L''(b)$ lies on any light-sheet emanating from a slice. Worse, there may not exist any choice of time slicing for which all of $L''(b)$ is covered by the light-sheets from the slices. The stronger bound $H(b) \subset L'(b)$ (which however is weaker than the bound we examine here) does ensure these properties, and we intend to investigate it further in future work.

In summary, the requirement that the bulk have no more degrees of freedom than the boundary, combined with arguments of symmetry and simplicity, has led us to propose the upper bound

$$L(b) \supset H(b) \quad (3.2)$$

on the holographic domain, where $L(b) \equiv L^+(b) \cup L^-(b)$ and $L^+(b)$ ($L^-(b)$) is the union of all future-directed (past-directed) light-sheets of the time slices that foliate the boundary region b .

IV. PROOF THAT $L = C$

In the previous two sections we have argued that the holographic bulk domain H dual to the boundary region b must satisfy

$$C(b) \subset H(b) \subset L(b). \quad (4.1)$$

In this section, we will prove that

$$C(b) = L(b). \quad (4.2)$$

This implies that

$$C(b) = H(b) = L(b), \quad (4.3)$$

so the holographic dual H is completely determined by our assumptions.⁴

It is obvious that $L(b) \subset C(b)$. But the converse inclusion is nontrivial, particularly since $L(b)$ is constructed from two sets $L^\pm(b)$ that depend on the slicing of b , whereas $C(b)$ is slicing independent. By Eq. (4.3), $C(b) \subset L(b)$ is required for the consistency of the arguments we have put forward in the previous section. Thus, our proof also serves as a nontrivial consistency check. It is the main technical result of this paper.

We begin by stating our assumptions and definitions. Let B be a manifold with boundary, and let $b \subset \partial B$ be a timelike embedded submanifold in B . We will require that b is globally hyperbolic when considered as a spacetime on its own. We also assume that B has the property that $J^\pm(P)$ is closed for every $P \in B$. Note that we are not assuming inextendibility of either b or B , nor global hyperbolicity of B (from which the latter assumption would follow).

⁴For certain choices of b , H thus excludes bulk points that can be represented on b [19,20]; see, however, Sec. VI. Our result is consistent with the fact that only C is needed to compute CFT correlation functions in b [21].

In the application of our theorem to AdS/CFT, we take b to be a portion of the conformal boundary of AdS. In this case, the spacetime B should be taken as the union of the unphysical conformally rescaled AdS spacetime and its boundary. Since our theorem only relies on properties of the spacetime which are preserved by conformal transformations, we are free to construct the proof in this unphysical spacetime.

An additional assumption is that the causal relation between any two points in b computed according to the causal structure of the lower-dimensional spacetime is the same as that according to the causal structure of B itself. This is essential for a physical duality to hold: the only way to guarantee that causality is preserved on both sides of the duality is to make the causal structures compatible in this way. (In asymptotically AdS spacetimes, this assumption follows from a theorem of Gao and Wald [22]. However, this theorem relies on additional assumptions that we have no reason to make here.)

Let $\tau: b \rightarrow \mathbf{R}$ be an arbitrary time function⁵ on b such that the equal time slices are Cauchy surfaces of b . The existence of τ is guaranteed by theorem 8.3.14 of Ref. [23]. Let $K \subset \mathbf{R}$ be the image of τ . Note that K is an open interval. Let $s_t \subset b$ be the Cauchy surface consisting of points with time t , $s_t = \{p \in b | \tau(p) = t\}$. Note that each s_t is spacelike. We will not demand that τ extends to a time function on B .

For the purpose of this theorem, we shall define $\ell_t^+ \subset B$ ($\ell_t^- \subset B$) as the set of future-directed (past-directed) null geodesics which are orthogonal to s_t with no conjugate points between s_t and the endpoint. (ℓ_t^\pm is a light-sheet associated to s_t if it is initially nonexpanding away from s_t . For the case of interest, where b is a portion of the conformal boundary of AdS, this always holds in the *physical* spacetime; see the Appendix. In the proof we will not make use of the nonexpansion property and so will not demand that ℓ_t^\pm be a true light-sheet in the unphysical spacetime.) Let $L^\pm = \bigcup_t \ell_t^\pm$, and $L = L^+ \cap L^-$.

Theorem: $L = C$ for any choice of time function on b , where $C \subset B$ is the set of points P that lie on a causal curve that begins and ends on b .

Corollary: L does not depend on the choice of time function, even though L^+ and L^- do.

Proof: Trivially, $L \subset C$. It remains to be shown that $C \subset L$, for all time functions τ (i.e., foliations of b). We will show that $C \subset L^-$, and similar arguments show that $C \subset L^+$. Let $P \in C$ (and P is not contained in b).

Definition.—Let $K^+(P)$ be the subset of the real numbers consisting of all t such that $s_t \cap I^+(P)$ is not empty. That is, $\{s_t | t \in K^+(P)\}$ is the set of all time slices on b which contain endpoints of future-directed timelike curves from P .

⁵That is, τ is differentiable on b , and $\nabla^a \tau$ is a past-directed timelike vector field.

Lemma A.— $K^+(P)$ is nonempty.

Proof.—Since $P \in C$, there is a point $p \in b \cap J^+(P)$. All points in $b \cap I^+(p)$ are necessarily in $I^+(P)$ (corollary following theorem 8.1.2 of Ref. [23]). Since K is an open interval, there exists $t \in K$ with $t > \tau(p)$. Any inextendible timelike curve passing through p intersects s_t because s_t is a Cauchy surface, which shows that $s_t \cap I^+(p)$, and consequently $b \cap I^+(P)$, is nonempty. Therefore $K^+(P)$ is nonempty.

Lemma B.—For any $t \in K^+(P)$, the set $A_t = b \cap J^+(P) \cap J^-(s_t)$ is compact.

Proof.—Since $P \in C$, $b \cap J^-(P)$ is not empty. Let $q \in b \cap J^-(P)$. Let $j^+(q)$ be the causal future of q within the spacetime b , and define $j^-(s_t)$ similarly. Then $j^+(q) \cap j^-(s_t)$ is compact (theorem 8.3.12 of Ref. [23]). But because the causal relation between points in b is the same whether we treat them as events in the spacetime b or the spacetime B , we also have that $j^+(q) \cap j^-(s_t) = b \cap J^+(q) \cap J^-(s_t)$. Since this set is compact, it is closed as a subset of B . $J^+(P)$ is a closed subset of $J^+(q)$, and it follows that $b \cap J^+(P) \cap J^-(s_t) = J^+(P) \cap (b \cap J^+(q) \cap J^-(s_t))$ is a closed subset of $b \cap J^+(q) \cap J^-(s_t)$. Therefore $b \cap J^+(P) \cap J^-(s_t)$ is compact.

Definition.—Let $t_*(P) = \inf_{t \in K^+(P)} t$.

Lemma C.—For any $t \in K^+(P)$, $t_*(P) = \min_{p \in A_t} \tau(p)$. In particular, this means that there is a surface $s_{t_*(P)}$ and that $s_{t_*(P)} \cap J^+(P)$ is not empty.

Proof.—By Lemma B, A_t is compact and hence $\tau(p)$ attains a minimum value τ_{\min} . Consider a point $p \in s_{\tau_{\min}} \cap J^+(P)$ and an inextendible future-directed timelike curve γ in b starting at p . All points on γ other than p are necessarily in $b \cap I^+(P)$ (corollary following Theorem 8.1.2 of Ref. [23]), and γ necessarily intersects $s_{t'}$ for every $t' \in K$ with $t' > \tau_{\min}$ because $s_{t'}$ is a Cauchy surface. This implies that all such times t' are in $K^+(P)$, and hence $\tau_{\min} \geq t_*(P)$. If $\tau_{\min} > t_*(P)$, then there is some $t' \in K^+(P)$ with $\tau_{\min} > t'$. Then $t > t'$, and so $s_{t'} \cap A_t$ is not empty by the definition of A_t . But then $\tau_{\min} \leq t'$ by the definition of τ_{\min} , which contradicts $\tau_{\min} > t'$. Therefore we conclude that $\tau_{\min} = t_*(P)$.

Lemma D.— $K^+(P)$ is an open subset of the real numbers. This implies that $s_{t_*(P)} \cap I^+(P)$ is empty.

Proof.—Let γ_p be an inextendible timelike curve in b which passes through a point $p \in b \cap I^+(P)$. We can choose to parametrize γ_p by the time function τ , which means that $\gamma_p: K \rightarrow b$ is a continuous function satisfying $\tau(\gamma_p(t)) = t$. Since $I^+(P)$ is open, it follows that the inverse image $\gamma_p^{-1}[I^+(P)]$ is open in K , and therefore open in \mathbf{R} because K is an open interval. By applying $\tau \circ \gamma$, we see that $\gamma_p^{-1}[I^+(P)] \subset K^+(P)$. Then $\bigcup_{p \in b \cap I^+(P)} \gamma_p^{-1}[I^+(P)]$ is an open subset of \mathbf{R} equal to $K^+(P)$.

Proof of theorem (see Fig. 8).—Lemma C and Lemma D together demonstrate that there is a surface $s_{t_*(P)}$ such that $s_{t_*(P)} \cap I^+(P)$ is empty while $s_{t_*(P)} \cap J^+(P)$ is nonempty.




FIG. 8. Consider an arbitrary boundary region b (enclosed by the solid line), and a point P in the bulk region $C(b)$ (orthogonal to the page). The causal future of P , $J^+(P)$, intersected with b is shown hatched. Roughly, the strategy of the proof is to demonstrate that there exists a time slice $s_{t_*(P)}$ on the boundary that is tangent to the lower boundary of future of P in b . We show that $s_{t_*(P)}$ is the earliest time slice that has any intersection with the future of P , and that P lies on the past light-sheet of $s_{t_*(P)}$.

Let $p \in s_{t_*(P)} \cap J^+(P)$. By the corollary following Theorem 8.1.2 of Ref. [23], there is a null geodesic connecting p to P . Furthermore, since $s_{t_*(P)} \cap I^+(P)$ is empty, this null geodesic cannot be deformed to a timelike curve connecting $s_{t_*(P)}$ to P . Then by Theorem 9.3.10 of Ref. [23], this null geodesic must be orthogonal to $s_{t_*(P)}$ and have no conjugate points between P and p . This shows that $P \in \ell_{t_*(P)}^- \subset L^-$.

V. COVARIANT RENORMALIZATION GROUP FLOW

We repeatedly made use of the UV/IR connection [11] in motivating our constructions. Entropy bounds play a very important role in UV/IR: the bulk region within the IR cutoff must have entropy limited by the area in Planck units of a time slice on the cutoff surface. However, as we have stressed, this only holds true for very special time slices and the covariant entropy bound must be used in general to bound the bulk entropy. Holographic renormalization group flows [24–27] aim to refine UV/IR, but in all standard approaches manifest covariance is lost and the status of entropy bounds is unclear. Here we outline an approach which reproduces the standard results and remedies both of these problems. Our construction gives an improved bulk radial flow, however it does not address the open question of finding a precise field theory RG representation of the bulk flow.

Choose a time function τ on b . Then for any t_1, t_2 the set $b[t_1, t_2] = \{p \in b | t_1 < \tau(p) < t_2\}$ satisfies the conditions of our theorem and we can associate to it the region $H(b[t_1, t_2])$. Now introduce a cutoff time scale δ to the theory. In the bulk, we should remove the union (over t) of the sets $H(b[t - \delta/2, t + \delta/2])$ from $H(b)$. The remaining bulk region, H^δ , is the region described by the cutoff CFT. This prescription is similar to the construction of [28], where it was shown that the IR cutoff surface as normally defined can be reproduced using only causality.

Denote the boundary of H^δ by b^δ . The time function on b automatically induces a time function on b^δ .⁶ By construction, b^δ is the union of sets of the form $s_t^\delta = \ell_{t+\delta/2}^- \cap \ell_{t-\delta/2}^+$, which we can take as time slices on b^δ . The maximum proper time in b^δ between s_t^δ and $s_{t+\delta}^\delta$ will be of order the AdS time. More generally, we can let δ depend on t .

To change the cutoff from δ to δ' , we have two equivalent options: First, we can return to the true boundary b and repeat the construction with δ' in place of δ . Second, we can use the surface b^δ , together with its induced time function, as the starting point for the construction, with cutoff $\delta' - \delta$. The geometry is the same either way, because the light-sheets from b^δ are continuations of the light-sheets from b .

By the definition of the L , and since $H = L$, all time slices of b^δ are *normal*. That is, their orthogonal light rays are everywhere nonexpanding in the direction of the bulk RG flow. A generic time slicing on b^δ will *not* admit ingoing past- and future-directed light-sheets at all points on all time slices; only the slicing induced by the flow has this property. And generic hypersurfaces other than those induced by the flow may not admit any slicing with this property. Note that this property is inherited from the remarkable property of the conformal boundary described in the Appendix: any slicing of the conformal boundary is everywhere normal in the physical metric.

The fact that all time slices of b^δ are normal is highly nontrivial. It leads to two additional, attractive features that distinguish this geometric flow from, say, the flow along spacelike geodesics:

- (i) Entropy bounds hold for both the UV and the IR regions.
- (ii) These bounds guarantee that the effective number of degrees of freedom is nonincreasing along the flow.

Consider first the IR region. The covariant entropy bound guarantees that any future- or past-directed light-sheet from any slice s_t^δ has entropy less than the area of that slice, in Planck units. Since the area of the light-sheet is nonincreasing as it moves away from the boundary, the number of degrees of freedom is nonincreasing along the RG flow.

Now consider the UV region, i.e., follow the light-sheet from $s_t^{\delta'}$ backward to some cutoff $\delta < \delta'$ closer to the true boundary. The cross-sectional area will be larger on b^δ and we can consider the area difference $A - A'$. A generalized version of the covariant entropy bound [29] implies that the entropy on the *partial light-sheet* between A and A' is bounded by $(A - A')/4$. Note that this bound applies to both light-sheets bounding the wedge-shaped region between the slices $s_{t+(\delta'-\delta)/2}^\delta$ and $s_{t-(\delta'-\delta)/2}^\delta$ on b^δ and $s_t^{\delta'}$ on




FIG. 9 (color online). The covariant bulk RG flow presented here reproduces the standard bulk RG flow in certain coordinate systems. Here we illustrate the construction in global coordinates of anti-de Sitter space. For a given coordinate time cutoff δ , the union over t of the intersection surfaces $s_t^\delta = \ell_{t+\delta/2}^- \cap \ell_{t-\delta/2}^+$ form a timelike hypersurface b^δ in the bulk (left). The cross-sectional area of a given light-sheet will be greater on the surface b^δ than on $b^{\delta'}$ (right). The difference $(A - A')/4$ bounds the entropy on the red light-sheets going from b^δ to $b^{\delta'}$, meaning that the bound applies to the entire darkly shaded wedge between them. The lightly shaded region between the hypersurfaces $b^{\delta'}$ and b^δ is covered by such wedges.

$b^{\delta'}$, and hence to the entire information content of the regions that are integrated over in the RG flow.⁷

RG flows are normally defined in terms of proper distances, times, or energies, whereas the above construction is in terms of an arbitrary coordinate time. To define proper distances on the boundary, one has to choose a conformal frame, and the choice of time function in our construction is analogous. Heuristically, one can imagine choosing the conformal factor so that τ becomes proper time, which means that it is conjugate to energy.

We have emphasized the geometric aspects of the RG flow, i.e., the bulk side. The question of what precisely the removal of near-boundary regions corresponds to in the CFT remains subtle, and it will be interesting to revisit it in light of the geometric flow we have described. In particular, it would be nice to understand whether a duality holds for the wedge-shaped regions associated with thin boundary strips, and for the Rindler portion of the bulk dual to small causal diamonds on the boundary. We are currently studying this question [30]; see also [5,31,32].

As the example of the global boundary (Fig. 9) illustrates, the description of the bulk region $H(b[t - \delta/2, t + \delta/2])$, for small time intervals δ , cannot involve the full set

⁶This will be true for sufficiently well-behaved time functions t ; the precise conditions will be investigated elsewhere.

⁷In general, the area $A(t)$ of different slices on b^δ will not be independent of t . It seems likely that requiring this independence will lead to preferred choices of $\delta(t)$ and thus for the flow. Also, while the areas will be automatically finite at finite δ in the case of the global boundary, they will be infinite in general and must be regulated.

of CFT degrees of freedom, since approximately $\exp(A/4)$ CFT states correspond to bulk states localized to the interior of the bulk hypersurface b^δ . If a duality exists, the CFT degrees of freedom relevant for $H(b^\delta)$ will not involve certain nonlocal operators that occupy scales larger than δ .

VI. NONLOCAL OPERATORS AND THE BULK DUAL

We constructed a candidate for a bulk region $H(b)$ dual to b geometrically, using considerations of causality and holographic entropy bounds. Except for appealing to the UV/IR relation [11], we did not make use of detailed properties of the AdS/CFT duality or the boundary theory. In this section we will explore a different approach to this question, namely the use of nonlocal operators in the CFT to probe deeply into the bulk region. We will examine the relationships between this approach and the geometric one. Note that the operator approach is available only for choices of b where the construction of boundary duals to local bulk operators is known, so it is less general than our geometric construction.

The methods of the operator approach were recently discussed by Heemskerck, Marolf, and Polchinski [20], and we refer the reader to that work for further details. The problem is to identify a subset of operators $\{\phi(x)\}$ within the set of all operators in the CFT on b , indexed by a position x in a semiclassical spacetime of higher dimension, which can be identified as local bulk operators. An important point is that the definition of the local bulk operators depends on the background metric. Our analysis below pertains to the case where this metric is held fixed up to perturbative corrections in $1/N$, so this issue should not pose difficulties.

The CFT definition of $\phi(x)$ will involve nonlocal CFT operators known as precursors [19,33,34]. An (in principle) explicit construction is available in the case of global AdS, and the resulting operators make use of an entire Cauchy surface of the global boundary. For this reason, and also for simplicity, we will spend most of this section focused on the case of a short strip $b = \mathcal{S}$,

$$\mathcal{S} = (-\tau_0, \tau_0) \times S^d, \quad \text{with } \tau_0 \ll 1, \quad (6.1)$$

which we normally think of as being embedded in the global boundary (see Fig. 10), using the coordinates of Eq. (2.1). From the global point of view, the set of operators on the Cauchy surface $\tau = 0$ of the boundary is complete, and so in particular contains the operator $\phi(x)$ for every point x in the entire global bulk.

Now consider modifications of the CFT Hamiltonian \mathcal{H} outside b , and let us define \bar{H} as the largest bulk region such that operators $\phi(x)$ for $x \in \bar{H}$ can be represented in terms of CFT operators in b in an unambiguous way. There are some modifications of the CFT Hamiltonian which continue to yield a well-defined bulk Hamiltonian. For




FIG. 10 (color online). A cross section of anti-de Sitter space, showing a short strip region \mathcal{S} centered around $\tau = 0$ on the boundary, and the bulk region $S(\mathcal{S})$ spacelike separated from \mathcal{S} . A local operator at the origin of the bulk can be written in terms of local operators on the boundary smeared over the boundary region spacelike related to the origin, within the green wedges. This region is much larger than \mathcal{S} (red thick line), stretching from $\tau = -\pi/2$ to $\tau = +\pi/2$.

instance, we can insert a local source in the CFT whose effect in the bulk is to cause a particle to propagate causally inward from the boundary. As we discuss below, certain modifications of the CFT Hamiltonian which make use of nonlocal operators can lead to ambiguities in the bulk Hamiltonian. In the case where we allow only those modifications leading to well-defined bulk Hamiltonians, we identify a region S on which the operators $\phi(x)$ have unambiguous CFT representations in b . In the more general case of modifications which lead to an ambiguous bulk Hamiltonian, we will identify \bar{H} as the subset of S for which $\phi(x)$ still has an unambiguous CFT representation in b . We will note that \bar{H} is closely related to H .

A. The region S

The only known construction of $\phi(x)$ in the boundary theory consists of two steps and applies either in global AdS or the Poincaré patch. In the first step, one writes a local bulk operator $\phi(x)$ as a smeared local operator on the boundary:

$$\phi(x) = \int dy' K(x|y') \mathcal{O}(y') + O(1/N). \quad (6.2)$$

Here y' denotes a boundary coordinate, while x is the bulk coordinate.⁸ The smearing function $K(x|y)$ is not unique,

⁸The presence of a source will modify this equation in the appropriate way given by Green's identity.

but a convenient choice is nonzero if and only if y spacelike related to x [5,31,32]. One can think of $K(x|y)$ as providing the solution to the equation of motion for $\phi(x)$ given boundary data on the asymptotic boundary; this is a space-like analogue of the standard initial value problem. Note that for any choice of x in global AdS, the boundary support of $K(x|y)$ is larger than the region \mathcal{S} . The second step of the construction uses unitary evolution in the CFT to write all of the local CFT operators $O(y')$ appearing in Eq. (6.2) in terms of nonlocal operators defined at $\tau = 0$ (or some other time slice within the short strip \mathcal{S}).

Now let us identify a subset $S(\mathcal{S})$ of the bulk such that the above construction of the CFT operator $\phi(x)$ is independent of modifications of the boundary Hamiltonian outside b ,⁹ provided that these modifications continue to lead to a well-defined bulk Hamiltonian. In Ref. [20], it was shown that S consists of the bulk points that are neither in the causal future of the future boundary of \mathcal{S} , nor in the causal past of its past boundary. We will refer to such points somewhat loosely as being “spacelike related” to \mathcal{S} (see Fig. 10).

At first, this result may seem surprising, so it is worth reviewing the argument for it [20]. First, fix a fiducial CFT Hamiltonian \mathcal{H} over the entire global boundary, and follow the above procedure to construct an operator \mathcal{O} satisfying $\phi(x) = \mathcal{O}$, for some x spacelike related to \mathcal{S} (the special case $x = 0$ is pictured in Fig. 10). But there is another way we can produce the operator \mathcal{O} : We can use causal bulk evolution to write $\phi(x)$ in terms of $\phi(x')$ for points x' in the future of the x :

$$\phi(x) = \int_{\Sigma} dx' [\phi(x') \nabla G_{\text{adv}}(x|x') - \nabla \phi(x') G_{\text{adv}}(x|x')] + O(1/N), \quad (6.3)$$

where Σ consists of a bulk Cauchy surface and, possibly, a portion of the global boundary. We can use Eq. (6.2) to write the $\phi(x')$ appearing here in terms of operators on the boundary. By evolving $\phi(x)$ sufficiently far forward into the future, the y' support of $K(x'|y')$ will lie in the region $\tau > 0$ for every x' appearing in Eq. (6.3). Now we evolve this new smeared operator back to $\tau = 0$ to obtain a second operator \mathcal{O}' . However, $\mathcal{O}' = \phi(x) = \mathcal{O}$, and so these two procedures actually give the same answer. Now suppose we modify the CFT Hamiltonian from our fiducial choice \mathcal{H} to $\tilde{\mathcal{H}}$, and we stipulate that $\tilde{\mathcal{H}}$ only differs from \mathcal{H} for $\tau < 0$. We can repeat the procedure to compute new operators $\tilde{\mathcal{O}}$ and $\tilde{\mathcal{O}}'$, which are equal to each other and to $\phi(x)$. The claim is that, since the calculations of \mathcal{O}' and $\tilde{\mathcal{O}}'$ refer only to the $\tau > 0$ region of the bulk and boundary, and $\mathcal{H} = \tilde{\mathcal{H}}$ in that region, the computations are identical and so manifestly we have

⁹We impose the boundary condition that the boundary state at $\tau = 0$ remain unchanged.

$\mathcal{O}' = \tilde{\mathcal{O}}'$. Therefore $\tilde{\mathcal{O}} = \mathcal{O}$, and the change in fiducial Hamiltonian did not change the operator assignment. An analogous argument can be made for modifications to the CFT Hamiltonian for $\tau > 0$.

B. The region \tilde{H}

The above argument relied crucially on the existence of a well-defined bulk Hamiltonian. However, there are reasonable modifications of the CFT Hamiltonian \mathcal{H} for which this will not be the case (see Fig. 11). At the time τ_0 on the boundary, let us add to \mathcal{H} a source for the nonlocal CFT operator \mathcal{O} dual to a local bulk operator at the origin, $\phi(\tau = 0, \rho = 0)$

$$\mathcal{H} \rightarrow \mathcal{H} + J\delta(\tau - \tau_0)\mathcal{O} \quad (6.4)$$

We are working perturbatively, so the operator \mathcal{O} is the one constructed using the above method and the unmodified Hamiltonian \mathcal{H} . In the bulk, this source acts completely locally as a source for ϕ at the origin at $\tau = 0$. Note that this bulk point is spacelike related to the boundary slice τ_0 .

Now let us compute the expectation value of the bulk operator ϕ at the origin, at some infinitesimal time ϵ after the source acts. This operator can be constructed by the usual methods, but those methods require a knowledge of the bulk Hamiltonian in the region $S(\mathcal{S})$. This Hamiltonian




FIG. 11 (color online). According to the Hamiltonian on the boundary strip \mathcal{S} , no source acts at the origin in the bulk, so the expectation value of ϕ vanishes everywhere. At the time τ_0 outside the strip, a source term for the nonlocal boundary operator dual to $\phi(0, 0)$ can be added to the boundary Hamiltonian. This causes the expectation value of ϕ to be nonzero in the future of $(0, 0)$, in contradiction with the earlier conclusion about the same bulk points. Thus, unless we possess information about the exterior of \mathcal{S} on the boundary guaranteeing that such operators do not act, the bulk interpretation of regions outside $\tilde{H}(\mathcal{S}) = C(\mathcal{S})$ is potentially ambiguous.

is ambiguous: from the viewpoint of the strip \mathcal{S} , the source does not act, since it acts in the CFT only at τ_0 . Then the bulk evolution should be computed from the usual bulk Hamiltonian without source, and we obtain $\langle\phi(\epsilon, 0)\rangle = 0$. From the viewpoint of the CFT at the time τ_0 , the bulk region near $(0,0)$ contains a source. Then the bulk evolution should take this source into account and we obtain $\langle\phi(\epsilon, 0)\rangle \neq 0$. Thus, there is no unique assignment of a bulk field value at $(\epsilon, 0)$. In this sense, $(\epsilon, 0)$ should not be considered a bulk point dual to the strip.

Let us now consider a general boundary region b and construct a bulk region \bar{H} such that the interpretation of what happens in \bar{H} is unambiguous (Fig. 12). The ambiguities we discussed arise from inserting nonlocal CFT operators into the Hamiltonian on Cauchy surfaces of the boundary which do not intersect the region b . A modification of the CFT Hamiltonian on such a Cauchy surface σ can lead to an ambiguous bulk Hamiltonian in the region $S(\sigma)$ spacelike related to σ . Thus we find that \bar{H} is the complement of the union of all $S(\sigma)$, where σ is any Cauchy surface for the global boundary which lies in the complement of b .

The region \bar{H} is closely related to $C(b)$ and hence to $H(b)$. It is easy to see that $H = C \subset \bar{H}$. The study of a number of examples suggests that $H \neq \bar{H}$ if and only if an event horizon is present in the bulk. It would be nice to study \bar{H} and its relation to H further. The discussion in the following section may be relevant.

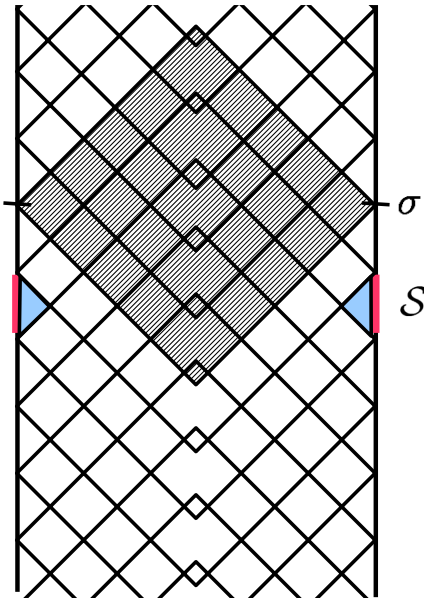


FIG. 12 (color online). The shaded region shows bulk points spacelike related to a global boundary Cauchy surface σ . The union of all such sets over the collection of boundary Cauchy surfaces which do not intersect \mathcal{S} has an ambiguous bulk interpretation when the boundary Hamiltonian is allowed to vary outside of \mathcal{S} . The unambiguous region, $\bar{H}(\mathcal{S})$, is the complement of this union. In this example, we see that $\bar{H}(\mathcal{S}) = C(\mathcal{S})$.

VII. QUANTUM GRAVITY BEHIND EVENT HORIZONS

In this section we discuss an issue that is somewhat orthogonal to the main subject of this paper: the degree to which the reconstruction of bulk regions, perturbatively in $1/N$, allows us to claim that AdS/CFT provides a full quantum gravity theory for regions behind event horizons, such as the interior of a black hole, or cosmological regions. We will construct an experiment behind the horizon whose outcome is known but not captured by such methods.

The CFT provides a full quantum gravity theory for observers near the boundary.¹⁰ It completely settles the issue of whether the formation and evaporation of a black hole is a unitary process. It is crucial for this argument that the time evolution is carried out on the boundary, where it is manifestly unitary; the duality is used only at early and late times in order to recover the in and out states in the bulk [35–38].

To what extent can we regard the CFT *also* as a full quantum gravity theory for an observer falling into the black hole? Perhaps, by reordering the degrees of freedom, one could interpret the CFT as providing a nonperturbative definition of quantum gravity for the infalling observer? This would require that the bulk dual region is ambiguous, at the nonperturbative level. This may be the case, but the dictionary that would provide this definition is not known at the required level of precision.

The black hole interior is clearly encoded in the CFT if one makes use of bulk equations of motion to evolve the infalling data back out of the black hole and to the boundary. But in the same approximation, we can also generate a xeroxing paradox [39]: at the semiclassical level, there is no manifest obstruction to evolving to global bulk slices that contain both the black hole interior and the Hawking cloud.

Perhaps we should restrict the bulk evolution by hand to the causal patch of an infalling observer, and settle for this approximate description of the black hole interior? The finiteness of entropy bounds inside black holes imply that there cannot be exact observables associated with the infalling observer at late times. Thus, the description of the infalling observer should be less precise than that of the observer near the boundary, who has access to exact observables. Perhaps the need to use bulk evolution is simply a reflection of this intrinsic limitation?

In fact, however, it is clear that the infalling observer requires a theory that goes beyond semiclassical bulk evolution. This can be seen from the following thought-

¹⁰Note that all observers that remain outside the black hole have the same causal diamond, consisting of the exterior of the black hole. The covariant RG flow we described in Sec. V can be thought of as moving the observers deeper into the bulk, but note that the flow never enters the black hole.

experiment. Consider an infalling observer who performs scattering experiments inside a Schwarzschild-AdS black hole of radius R , after crossing the event horizon. The scattering occurs at high enough energy to produce a (smaller) black hole of mass m , which then evaporates. We are free to choose parameters so that the scattering effectively takes place in flat space:

$$R_{\text{AdS}} \gg R \gg f(m) \gg 1, \quad (7.1)$$

where $f(m)$ is at least the evaporation time scale of the black hole (m^d in d spatial bulk dimensions). For sufficiently large but finite $R/f(m)$, the infalling observer can confirm the unitarity of black hole formation and evaporation to any required precision. But this conflicts with the result that would be obtained from the semiclassical analysis (Hawking's calculation).

We conclude that the proper description of an infalling observer requires a quantum gravity theory beyond the semiclassical approximation. The CFT on the boundary does not provide us with that theory, since its application to the infalling observer relies on semiclassical bulk evolution. No other way of relating local operators inside the black hole to boundary operators is known; therefore, we cannot replace this bulk evolution by boundary evolution, as we did in the case of scattering experiments performed by a near-boundary observer.

Of course, in the limit used in the thought experiment, one could imagine “cutting out” the spacetime region of size $f(m)$ that contains the scattering experiment. One could embed this region in an asymptotically AdS spacetime with $R_{\text{AdS}} \gg f(m)$ and use the flat space S-matrix, which can be computed using AdS/CFT [35–38]. In fact, this argument is what gives us confidence that the process is indeed unitary. But how would this prescription generalize? A theory of the infalling observer that relied on this type of cutting and pasting would not be applicable to the highly dynamical regions deeper inside the black hole, nor to the spacelike singularity, which cannot be so transplanted.

An exact version of this cut-and-paste process is available if the black hole is formed by sending in a spherical null shell from the boundary of AdS [40]. By causality, the bulk region in the past of the shell is the same no matter whether we decide to send in the shell or not. If we do send in the shell, then this region includes a portion w of the black hole interior. If we do not, then the same bulk region w is dual to operators on the boundary, which can be evolved to nonlocal operators W on a single global boundary time slice preceding the insertion of the shell. This illustrates that cut-and-paste is well defined precisely in the limit where it yields no information about regions that could *only* exist in the interior of black holes, such as singularities and a highly dynamical geometry.

The absence (so far) of an intrinsic bulk theory at the nonperturbative level appears to impose crucial

limitations on our ability to describe black hole interiors and cosmological regions¹¹ via AdS/CFT, beyond what follows from the approximate methods that were already available to generate bulk evolution. (See, however, Refs. [44–47] for interesting approaches to this problem.)

ACKNOWLEDGMENTS

We thank N. Arkani-Hamed, T. Banks, B. Freivogel, J. Kaplan, H. Liu, J. Maldacena, D. Marolf, J. Polchinski, E. Rabinovici, S. Shenker, A. Strominger, L. Susskind, and E. Witten for discussions. This work was supported by the Berkeley Center for Theoretical Physics, by the National Science Foundation (Awards No. 0855653 and No. 0756174), by fqxi Grant No. RFP3-1004, and by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The work of V.R. is supported by an NSF Graduate Fellowship. The work of S.L. is supported by a John A. McCone Postdoctoral Fellowship.

APPENDIX: LIGHT-SHEETS FROM THE CONFORMAL BOUNDARY OF ADS

In this Appendix we show that the light-rays orthogonal to any spacelike slice of the conformal boundary of AdS have nonpositive expansion into the bulk in the physical metric, and thus generate light-sheets. This is an important property of light-sheets in asymptotically AdS spacetimes. It guarantees that the maximum entropy of the bulk holographic domain H is nonincreasing under the covariant RG flow described in Sec. V.

In Poincaré coordinates, the AdS metric is

$$ds^2 = \frac{1}{z^2} (dz^2 + \eta_{\mu\nu} dx^\mu dx^\nu), \quad (\text{A1})$$

where $\eta_{\mu\nu}$ is the metric for d dimensional Minkowski space. Consider the conformally rescaled metric

$$d\tilde{s}^2 = dz^2 + \eta_{\mu\nu} dx^\mu dx^\nu, \quad (\text{A2})$$

which is $d + 1$ dimensional Minkowski space. Take some small region of the boundary and let $\tilde{\theta}$ be the expansion of some congruence of infinitesimally neighboring light rays in this space. That is

$$\tilde{\theta} = \frac{d \log \tilde{\mathcal{A}}}{d\tilde{\lambda}}, \quad (\text{A3})$$

¹¹A similar conclusion [41] applies to the interior of other event horizons, such as a Friedman-Robertson-Walker universe [42]. In this case there are two natural choices of conformal frame on the boundary [43]: one in which the CFT is well behaved, and another in which the coefficient of a relevant operator diverges in finite time. It is tempting to interpret this violent behavior in terms of an infalling observer hitting the big crunch singularity behind the horizon, but it can be understood more simply as the arrival of the bulk domain wall on the boundary.

where $\tilde{\mathcal{A}}$ is the infinitesimal area spanned by the light rays and $\tilde{\lambda}$ is the affine parameter. In order for the null geodesics to remain affinely parametrized after the conformal transformation, the affine parameter must transform as (Appendix D of [23])

$$\frac{d\tilde{\lambda}}{d\lambda} = cz^2, \quad (\text{A4})$$

where c is a constant and z^2 is the conformal factor. The area will transform as

$$\mathcal{A} = \frac{\tilde{\mathcal{A}}}{z^{d-1}}. \quad (\text{A5})$$

Using (A3)–(A5) we find that the expansion in the AdS spacetime (A1) is

$$\theta = cz^2\tilde{\theta} - c(d-1)z\tilde{k}^z, \quad (\text{A6})$$

where

$$\tilde{k}^z = \frac{dz}{d\tilde{\lambda}}.$$

Since both $\tilde{\theta}$ and \tilde{k}^z are defined in $d+1$ dimensional Minkowski space using a congruence of null rays orthogonal to a spacelike (and hence nowhere null) foliation, they are finite. Thus, at the boundary ($z=0$), we have $\theta=0$.

-
- [1] J. Maldacena, *Adv. Theor. Math. Phys.* **2**, 231 (1998).
- [2] E. Witten, *Adv. Theor. Math. Phys.* **2**, 253 (1998).
- [3] T. Banks, M.R. Douglas, G.T. Horowitz, and E.J. Martinec, [arXiv:hep-th/9808016](https://arxiv.org/abs/hep-th/9808016).
- [4] V. Balasubramanian, P. Kraus, A.E. Lawrence, and S.P. Trivedi, *Phys. Rev. D* **59**, 104021 (1999).
- [5] A. Hamilton, D.N. Kabat, G. Lifschytz, and D.A. Lowe, *Phys. Rev. D* **74**, 066009 (2006).
- [6] D. Harlow and D. Stanford, [arXiv:1104.2621](https://arxiv.org/abs/1104.2621).
- [7] S. Ryu and T. Takayanagi, *Phys. Rev. Lett.* **96**, 181602 (2006).
- [8] T. Takayanagi, *Phys. Rev. Lett.* **107**, 101602 (2011).
- [9] M. Fujita, T. Takayanagi, and E. Tonni, *J. High Energy Phys.* **11** (2011) 043.
- [10] V.E. Hubeny, M. Rangamani, and T. Takayanagi, *J. High Energy Phys.* **07** (2007) 062.
- [11] L. Susskind and E. Witten, [arXiv:hep-th/9805114](https://arxiv.org/abs/hep-th/9805114).
- [12] R. Bousso and L. Randall, *J. High Energy Phys.* **04** (2002) 057.
- [13] V.E. Hubeny, *J. High Energy Phys.* **07** (2012) 093.
- [14] G. 't Hooft, [arXiv:gr-qc/9310026](https://arxiv.org/abs/gr-qc/9310026).
- [15] L. Susskind, *J. Math. Phys. (N.Y.)* **36**, 6377 (1995).
- [16] R. Bousso, *J. High Energy Phys.* **07** (1999) 004.
- [17] R. Bousso, *J. High Energy Phys.* **06** (1999) 028.
- [18] R. Bousso, *Rev. Mod. Phys.* **74**, 825 (2002).
- [19] J. Polchinski, L. Susskind, and N. Toumbas, *Phys. Rev. D* **60**, 084006 (1999).
- [20] I. Heemskerck, D. Marolf, and J. Polchinski, [arXiv:1201.3664](https://arxiv.org/abs/1201.3664).
- [21] D. Marolf, *J. High Energy Phys.* **05** (2005) 042.
- [22] S. Gao and R.M. Wald, *Classical Quantum Gravity* **17**, 4999 (2000).
- [23] R.M. Wald, *General Relativity* (University of Chicago Press, Chicago, 1984).
- [24] V. Balasubramanian and P. Kraus, *Phys. Rev. Lett.* **83**, 3605 (1999).
- [25] J. de Boer, E.P. Verlinde, and H.L. Verlinde, *J. High Energy Phys.* **08** (2000) 003.
- [26] T. Faulkner, H. Liu, and M. Rangamani, *J. High Energy Phys.* **08** (2011) 051.
- [27] I. Heemskerck and J. Polchinski, *J. High Energy Phys.* **06** (2011) 031.
- [28] R. Bousso, *Phys. Rev. D* **79**, 123524 (2009).
- [29] E.E. Flanagan, D. Marolf, and R.M. Wald, *Phys. Rev. D* **62**, 084035 (2000).
- [30] R. Bousso, B. Freivogel, S. Leichenauer, V. Rosenhaus, and C. Zukowski (to be published).
- [31] A. Hamilton, D.N. Kabat, G. Lifschytz, and D.A. Lowe, *Phys. Rev. D* **73**, 086003 (2006).
- [32] D. Kabat, G. Lifschytz, and D.A. Lowe, *Phys. Rev. D* **83**, 106009 (2011).
- [33] L. Susskind and N. Toumbas, *Phys. Rev. D* **61**, 044001 (2000).
- [34] B. Freivogel, S.B. Giddings, and M. Lippert, *Phys. Rev. D* **66**, 106002 (2002).
- [35] L. Susskind, [arXiv:hep-th/9901079](https://arxiv.org/abs/hep-th/9901079).
- [36] J. Polchinski, [arXiv:hep-th/9901076](https://arxiv.org/abs/hep-th/9901076).
- [37] S.B. Giddings, *Phys. Rev. Lett.* **83**, 2707 (1999).
- [38] A. Fitzpatrick and J. Kaplan, [arXiv:1104.2597](https://arxiv.org/abs/1104.2597).
- [39] L. Susskind and L. Thorlacius, *Phys. Rev. D* **49**, 966 (1994).
- [40] V.E. Hubeny, *Int. J. Mod. Phys. D* **12**, 1693 (2003).
- [41] J. Garriga, *Prog. Theor. Phys. Suppl.* **190**, 261 (2011).
- [42] J. Maldacena, [arXiv:1012.0274](https://arxiv.org/abs/1012.0274).
- [43] J. Barbon and E. Rabinovici, *J. High Energy Phys.* **04** (2011) 044.
- [44] P. Kraus, H. Ooguri, and S. Shenker, *Phys. Rev. D* **67**, 124022 (2003).
- [45] L. Fidkowski, V. Hubeny, M. Kleban, and S. Shenker, *J. High Energy Phys.* **02** (2004) 014.
- [46] G. Festuccia and H. Liu, *J. High Energy Phys.* **04** (2006) 044.
- [47] G. Horowitz, A. Lawrence, and E. Silverstein, *J. High Energy Phys.* **07** (2009) 057.