

RECEIVED: May 21, 2019

REVISED: June 28, 2019

ACCEPTED: July 11, 2019

PUBLISHED: July 25, 2019

Towards a bit threads derivation of holographic entanglement of purification

Ning Bao,^{a,b,c} Aidan Chatwin-Davies,^d Jason Pollack^e and Grant N. Remmen^{a,b}

^aCenter for Theoretical Physics and Department of Physics, University of California, Berkeley, CA 94720, U.S.A.

^bLawrence Berkeley National Laboratory, Berkeley, CA 94720, U.S.A.

^cComputational Science Initiative, Brookhaven National Lab, Upton, NY 11973, U.S.A.

^dKU Leuven, Institute for Theoretical Physics, Celestijnenlaan 200D B-3001 Leuven, Belgium

^eDepartment of Physics and Astronomy, University of British Columbia, Vancouver, BC V6T 1Z1, Canada

E-mail: ningbao75@gmail.com, aidan.chatwindavies@kuleuven.be, jpollack@phas.ubc.ca, grant.remmen@berkeley.edu

ABSTRACT: We apply the bit thread formulation of holographic entanglement entropy to reduced states describing only the geometry contained within an entanglement wedge. We argue that a certain optimized bit thread configuration, which we construct, gives a purification of the reduced state to a full holographic state obeying a precise set of conditional mutual information relations. When this purification exists, we establish, under certain assumptions, the conjectured $E_P = E_W$ relation equating the entanglement of purification with the area of the minimal cross section partitioning the bulk entanglement wedge. Along the way, we comment on minimal purifications of holographic states, geometric purifications, and black hole geometries.

KEYWORDS: AdS-CFT Correspondence, Black Holes, Gauge-gravity correspondence

ARXIV EPRINT: [1905.04317](https://arxiv.org/abs/1905.04317)

Contents

1	Introduction	1
2	Background	2
2.1	Entanglement of purification	2
2.2	Bit threads	4
3	Upper bound on E_P	4
3.1	Bit threads as tags of purifying degrees of freedom	5
3.2	A flux-maximizing thread configuration	6
3.3	Factorization and derivation of the upper bound	11
4	Lower bound on E_P	14
5	Discussion	17
A	Holographic purification of minimal dimension	18

1 Introduction

The relation between entanglement and (holographic) spacetime [1–3] is a powerful connection that has been extended through many interesting developments over the past decade. Perhaps the most elegant of these extensions is the Ryu-Takayanagi (RT) formula,

$$S[\rho] = \frac{\mathcal{A}}{4G_N\hbar}, \tag{1.1}$$

relating the entanglement entropy $S[\rho] = -\text{Tr} \rho \log \rho$ of the reduced density matrix ρ of a holographic large- N CFT in a boundary region to the area \mathcal{A} of the minimal bulk surface homologous to that region [4–6]. This result, among others, has led to a concrete instantiation of the emergence of spacetime from entanglement in the context of AdS/CFT [7–10].

Recently, the connection between boundary information-theoretic quantities and bulk areas has been extended to a conjectured “ $E_P = E_W$ ” relationship between the area of bulk-anchored minimal surfaces and the entanglement of purification [11], as in refs. [12, 13]. The entanglement of purification is an entanglement measure that characterizes the degree of entanglement between two subsystems of a generically mixed state. It is particularly useful when only partial information about a state is known. As we review below, the entanglement of purification is a bound on the entanglement between subsystems in a given purification; purifications that saturate this bound comprise a class of optimal purifications or completions of the state in question. The $E_P = E_W$ conjecture has been further extended, studied, and generalized in recent work [14–22].

In this paper, we will seek to make progress towards proving the $E_P = E_W$ conjecture using technology from the bit threads program as described in ref. [23]. This program, which in its original formalism [23] is formally equivalent to the RT prescription, provides useful tools and visualizations for understanding the structure of entanglement in holographic states.¹ In essence, the bit threads formulation of the RT formula replaces minimization of areas of surfaces with maximizations of constrained flows of vector fields through surfaces. These vector fields are envisioned as providing, at least heuristically, connections between Bell pairs of maximally-entangled qubits in the boundary state.

The usual conceptual flow of the bit threads program is to pass from a given state to a bit thread geometry describing (aspects of) the entanglement structure of the state. In this paper — and in the context of a holographic correspondence between boundary states and bulk geometries — we wish to invert this logic: we conjecture that, under certain circumstances discussed below, the existence of a bit thread geometry satisfying particular geometric constraints *implies* the existence of a state with subsystem entanglements satisfying the analogues of these constraints. Given this conjecture, we can replace the (extremely hard) task of specifying a particular holographic state directly in the boundary CFT with a more tractable geometric construction on the level of a classical bulk geometry. Because we expect that geometrical information is encoded redundantly in the boundary theory in the manner of an error-correcting code (see, e.g., ref. [25]), we do not expect this construction to uniquely specify a quantum-mechanical state. Nevertheless, we show in this paper that it provides enough information to compute the information-theoretic quantity dual to the cross-sectional area E_W and hence establish that $E_P = E_W$.

The organization of this paper is as follows. In section 2, we review some facts about bit threads and the entanglement of purification. In section 3, we argue that the entanglement of purification is bounded from above by the area of the entanglement wedge cross section, while in section 4 we show, under certain assumptions that we clarify, that the entanglement of purification is lower-bounded by the area of the entanglement wedge cross section, establishing the $E_P = E_W$ relation. Finally, we conclude in section 5, with a bonus related result on the minimal dimension of holographic purifications given in appendix A.

2 Background

We begin by introducing some definitions and briefly reviewing the concepts of entanglement of purification and bit threads.

2.1 Entanglement of purification

Let $\rho_{AB} \in \mathcal{L}(\mathcal{H}_{AB})$ be a state in the bipartite Hilbert space $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. The *entanglement of purification* of ρ_{AB} is

$$E_P(A : B) = \inf_{A'B'} S(AA'), \tag{2.1}$$

¹The bit threads formalism has further been extended to include multiflows and thread configurations [24].




Figure 1. Entanglement wedge for AB , bounded by the RT surfaces γ_{AB} . The entanglement wedge cross section Γ is illustrated by the red dashed line.

where the infimum is taken over all auxiliary Hilbert spaces $A'B'$ and all pure states $|\Psi\rangle_{AA'BB'}$ such that $\text{Tr}_{A'B'} |\Psi\rangle\langle\Psi| = \rho_{AB}$. The entanglement of purification is an entanglement measure that quantifies the minimum amount of correlation required in any purification of a mixed state, subject to the constraint that it must preserve the factorization structure between A and B ; this is particularly useful, for example, when only a mixed subset of the full pure state is known.

When \mathcal{H}_{AB} is a subfactor of the Hilbert space of a holographic CFT and ρ_{AB} is the reduced state on boundary subregions A and B of a state with a well-defined dual geometry, the entanglement of purification E_P has been conjectured to be equal to the area E_W of the entanglement wedge cross section Γ [12]. The entanglement wedge cross section Γ is a surface of minimal area anchored to the boundary of the entanglement wedge W_{AB} , such that Γ partitions W_{AB} into a region that is entirely adjacent to A and a region that is entirely adjacent to B ; see figure 1.

In the case of finite-dimensional Hilbert space, $\dim \mathcal{H}_{AB} = d_{AB} < \infty$, the infimum $E_P(A : B)$ can be achieved with auxiliary Hilbert spaces of dimensions $d_{A'} = d_{AB}$ and $d_{B'} = d_{AB}^2$ [11]. Note that while the Hilbert spaces of AdS/CFT are infinite-dimensional, we will later consider UV regularizations that render them finite-dimensional. Therefore, we can assume that the infimum is achievable in regulated holographic settings.

The state that achieves the entanglement of purification as $S(AA')$ is of course not unique. (For example, given a state $|\Psi\rangle$ that achieves $E_P(A : B)$, any other state of the form $U_{A'} \otimes U_{B'} |\Psi\rangle$ also achieves $E_P(A : B)$.) One should therefore think of entanglement of purification as picking out an equivalence class of “optimal” states on Hilbert spaces of possibly different dimensions that are all isometrically related. In appendix A, we show that if one holds a purification of holographic ρ_{AB} in \mathcal{H}_{CFT} that achieves $E_P(A : B)$, then one can compress the purifying state on $(AB)^c$ to a Hilbert space with $\log \dim \mathcal{H}_{A'B'} = S(\rho_{AB})$.

2.2 Bit threads

Bit threads were first introduced as a by-product of a reformulation of the RT formula in terms of flows by Freedman and Headrick in ref. [23].² Here, we introduce only the minimum results required for our purposes, leaving the details to ref. [23].

Let \mathcal{M} be an oriented Riemannian manifold with boundary and let $C > 0$. A *flow* is a vector field v such that $\nabla_\mu v^\mu = 0$ and $|v| \leq C$. Given a flow, one can of course compute the flow's flux through any sufficiently smooth codimension-one surface m :

$$\int_m v = \int_m \sqrt{h} n_\mu v^\mu, \tag{2.2}$$

where h is the induced metric on m and n is the normal to m .

The crux of Freedman and Headrick's reformulation of the RT formula is the Max-Flow/Min-Cut (MFMC) Theorem [28–30], which can be stated as follows. Let A be a boundary subregion of \mathcal{M} . Then

$$\max_v \int_A v = C \min_{m \sim A} \text{area}(m), \tag{2.3}$$

where $m \sim A$ denotes that m is homologous to A . MFMC states that for any flow on \mathcal{M} that attains the largest possible flux out of A , the value of this flux is equal to the largest transverse flow density C multiplied by the area of a minimal surface in \mathcal{M} that is homologous to A . In other words, the max flow necessarily saturates any bottleneck out of A in both magnitude and direction.

Letting \mathcal{M} be a spatial slice of a static holographic geometry and $C = 1/4G_N\hbar$, it should now not seem surprising that one can rewrite eq. (1.1) as

$$S(A) = \max_v \int_A v. \tag{2.4}$$

This was rigorously demonstrated in ref. [23].

Bit threads are the integral curves of a max flow with transverse density $|v|$. As a result of eq. (2.4), one can think of $S(A)$ as counting the largest number of bit threads leaving A that can pass through the RT surface for A . Due to the non-uniqueness of the max flow, one can also compute further entropic quantities by counting bit threads of appropriately chosen flows. We will comment further on how this is done in section 3 when we begin our calculations of this type.

3 Upper bound on E_P

Let ρ_{AB} be a holographic CFT state on the subregions A and B whose geometric dual is the entanglement wedge of AB . In this section, we establish an upper bound on its entanglement of purification, $E_P(A : B) \leq E_W(A : B)$, under an assumption regarding the relation between the entanglement structure of geometric purifications of the state and

²The formulation of the RT prescription in terms of flows can also be understood [26] from the perspective of calibrations defined on Riemannian manifolds [27].




Figure 2. Types of bit threads that can cross the RT surface, γ_{AB} , of AB for an arbitrary flow. The properties that distinguish different types of bit threads are where they begin and end (A , B , or $(AB)^c$), as well as how many times they intersect γ_{AB} and the entanglement wedge cross section Γ (red dashed line).

a bit thread flow in the entanglement wedge. To do so, it suffices to exhibit a purification for which $S(AA') = E_W(A : B)$. We will show that an arbitrary purification that is everywhere holographically dual to an asymptotically-AdS bulk geometry admits a factorization $|\Psi\rangle_{AA'BB'Y}$ for which $S(AA') = E_W(A : B)$. Such a purification is guaranteed to exist by the fact that geometric subregions of an AdS geometry (in this case the entanglement wedge) can always be extended to a full asymptotically-AdS bulk geometry.³ We now specify how to determine $S(AA')$ for this state and for a particular choice of A' , which we fix by using the bit thread formalism to specify various conditional mutual informations.

3.1 Bit threads as tags of purifying degrees of freedom

Consider two boundary subregions A and B , as well as the bit threads of some flow on a geometric purification of ρ_{AB} . The ways in which bit threads can intersect the RT surface of AB (which in general consists of multiple disconnected pieces) are sketched in figure 2.

Threads that start in AB and terminate in $(AB)^c$ are to be associated with the subfactor of $\mathcal{H}_{(AB)^c}$ that purifies ρ_{AB} . That is, $(AB)^c$ contains many extraneous degrees of freedom that are uncorrelated with ρ_{AB} and hence do not purify our reduced density matrix. This can be seen, for example, by looking at small boundary subregions $X \subset (AB)^c$ that are far away from $\partial(AB)$, as shown in figure 3. Such regions have zero mutual information

³In the case where ρ_{AB} resulted from tracing out $(AB)^c$ in a holographic state, then one trivial everywhere-geometric purification is the original state. More generally, on the level of the bulk geometry, a geometric purification corresponds to gluing the entanglement wedge of AB to other subregions in such a way that the entire geometry is asymptotically AdS. In arbitrary dimensions such a gluing construction may be nontrivial for general geometries, but it has been shown that gluing another copy of the entanglement wedge to itself along RT surfaces is always possible, forming a canonical purification [19, 31]; in AdS_3 , where the Weyl tensor vanishes identically, even more general gluing constructions should be fairly straightforward. Note that it may be necessary to introduce additional copies of the CFT if, for example, AB is already a full boundary and ρ_{AB} is a mixed state, such as in the case of a thermal single-sided black hole.




Figure 3. Boundary subregions whose reduced density matrix approximately factorizes.

with AB and the state factorizes⁴ in the large- N limit:⁵

$$\rho_{ABX} = \rho_{AB} \otimes \rho_X. \quad (3.1)$$

As such, we conclude that it is only a subset of the degrees of freedom, $\mathcal{H}_{(AB)'} \subset \mathcal{H}_{(AB)^c}$, that purify ρ_{AB} ; bit threads are to be associated with this smaller set of boundary degrees of freedom. Although we have chosen notation here to suggest a division $\mathcal{H}_{(AB)'} = \mathcal{H}_{A'} \otimes \mathcal{H}_{B'}$, so far we have not actually specified such a division, nor in fact required that $\mathcal{H}_{(AB)'}$ even factorize. In particular, $\mathcal{H}_{(AB)'}$ need not have any geometric locality, beyond simply that it is supported in $(AB)^c$. This is a reflection of the freedom in choosing the flow, i.e., the fact that the location where the threads terminate in $(AB)^c$ can be slid around freely anywhere in $(AB)^c$, and so they are not associated with literal, local “boundary qubits.” For example, the threads could be evenly spaced out, or we could bunch them up together about a particular location on the boundary, etc.

3.2 A flux-maximizing thread configuration

In order to fix the division of $(AB)'$ into A' and B' , we now construct a specific configuration of bit threads that simultaneously maximizes the number of threads crossing the RT

⁴In field theory, the fact that the vacuum state is entangled at all scales implies that even local operators well outside the entanglement wedge might change the reduced state ρ_{AB} . A finite-energy particle excitation at a particular spacetime location in fact has (exponentially small) support on all of space. So strictly speaking, we should throughout replace any operator acting outside the entanglement wedge with a different operator, with cutoff-scale energy, that genuinely does not change the reduced state of the entanglement wedge. Or, perhaps preferably, we could instead not demand that the new reduced state be exactly the same as the old, but only the same up to exponentially small differences. That is, whenever we factorize the Hilbert space of a CFT, we can either replace statements of factorization with appropriate statements involving von Neumann algebras (and/or bulk reconstruction) or edge modes, or we can embrace the existence of the cutoff and work with a latticized theory that lacks these issues. In the remainder of this paper, we take the latter, more pedestrian approach.

⁵It should be noted that (even beyond the caveats just noted in footnote 4) this factorization is approximate, to leading order in $1/N$, since bit threads are equivalent to the RT formula, which only computes the entropies to leading order in $1/N$. We are helped, however, by the fact that the trace distance between the exact state and the factorized state discussed here will be small, so Fannes’ inequality [32] guarantees that the difference in the entanglement entropies of subsystems of the two states will also be correspondingly small.

surface of AB , which we denote by γ_{AB} , and the number crossing the entanglement wedge cross section Γ . Note, however, that such a collection of threads cannot be the flow lines of an everywhere-continuous and -divergenceless flow with bounded norm. According to MFMC (2.3), if a flow v indeed maximizes the flux out of AB , then v must be normal to γ_{AB} and have $|v| = 1/4G_N\hbar$ everywhere on γ_{AB} . But then, the flux through Γ cannot be strictly equal to $|\Gamma|/4G_N\hbar$ since, on the codimension-three surface where Γ intersects γ_{AB} , v is perpendicular to the normal of Γ .

Therefore, we must necessarily relax the definition of bit threads as integral curves of a flow. We will simply take bit threads to be boundary-anchored one-dimensional objects whose density is at most $1/4G_N\hbar$, defining thread density, as in ref. [14], as the length of threads within some small neighborhood divided by the volume of that neighborhood. This generalization was proposed in ref. [24], in which a collection of such threads was called a “thread configuration.” Although the authors of ref. [24] further relaxed the requirement that bit threads be oriented, we will find it helpful to still think of our bit threads as having an orientation.

The use of this more general notion of thread configurations is very mild for our purposes, since ultimately we will only be concerned with computing entropic quantities by counting bit threads, as opposed to exploiting MFMC to identify the bottleneck of a flow. In other words, we will always assume that the relevant extremal surfaces are known, and we will only need to count threads that cross these surfaces. Moreover, the only place where the associated flow becomes multivalued is on the (codimension-three) surface $\Gamma \cap \gamma_{AB}$, so with bounded flux density, the region of this ambiguity contributes only a measure-zero fraction of the flux through Γ or γ_{AB} .

With this generalization in mind, we can now describe how to construct a thread configuration that saturates the number of threads crossing γ_{AB} and Γ . Let us suppose that a UV cutoff near the boundary has been imposed, so that the areas of boundary-anchored surfaces are finite and well defined. First, begin by letting $S(A) = |\gamma_A|/4G_N\hbar$ thread fragments emanate from A and $S(B) = |\gamma_B|/4G_N\hbar$ thread fragments flow into B , where γ_A and γ_B denote the RT surfaces of A and B individually and where the thread fragments run from the boundary to γ_A and γ_B for now. These are the largest numbers of threads that we can pipe from A to the bulk and to B from the bulk without violating the density bound anywhere. The reader will note that we have suggestively labeled the numbers of threads using the same symbol as entropy, and we will proceed to work with “entropies” and “mutual informations,” but, as we discuss in the next subsection, *a priori* these quantities are purely geometric and need not be associated with the entanglement entropy of reduced density matrices constructed from a pure state.

One of our goals is to saturate the number of threads that cross γ_{AB} , so $S(AB) = |\gamma_{AB}|/4G_N\hbar$ of the fragments will have to cross γ_{AB} and leave the entanglement wedge. Therefore, take $[S(A) + S(B) - S(AB)]/2$ of the thread fragments from A and the same number from B (i.e., the leftover fragments) and join them by having them cross Γ . At this stage, there are thus $[S(A) + S(B) - S(AB)]/2 = I(A : B)/2$ bit threads crossing Γ .

Next, since $E_W(A : B) \geq I(A : B)/2$ (as shown in refs. [12, 23]), we may still need to pipe some of the thread fragments through Γ . This is just a matter of ensuring that

$E_W(A : B) - I(A : B)/2$ of the remaining $S(AB)$ fragments cross Γ before going on to intersect γ_{AB} , after which the threads are sent to the boundary where they terminate in $(AB)^c$. In particular, we can see that there are enough remaining threads to saturate Γ as follows:

$$\begin{aligned}
 S(AB) + I(A : B)/2 &= \frac{1}{2} [S(A) + S(B) + S(AB)] \\
 &\geq \max\{S(A), S(B)\} \\
 &\geq E_W(A : B).
 \end{aligned}
 \tag{3.2}$$

In going to the second line we used the Araki-Lieb inequality, and to go to the third line we used the inequality $E_W(A : B) \leq \min\{S(A), S(B)\}$ [12]. The result is a thread configuration with $S(AB)$ threads crossing γ_{AB} and $E_W(A : B)$ threads crossing Γ , thus saturating both surfaces as desired (see figure 4c).

Alternatively, we can obtain our γ_{AB} - and Γ -saturating thread configuration directly from the thread configuration constructed in section 5.2.2 of ref. [14], via either cutting-and-gluing or combing. There, Agón et al. construct a configuration of threads that saturates the number of threads crossing γ_{AB} , but not Γ (see their figure 11). From their configuration, one can arrive at a new configuration that saturates both quantities by cutting threads on either side of Γ and rejoining them with neighboring threads inside the entanglement wedge to pipe $E_W(A : B)$ threads through Γ ; this is illustrated in figure 4a-b. This is equivalent to “combing” the thread configuration by dragging the threads that intersect γ_{AB} in opposite directions along the interior and exterior disconnected components of γ_{AB} , as illustrated in figure 4d. We refer the reader to the caption of figure 4 for more details.

Agón et al.’s construction applies specifically to the case where AB is a proper subset of a single boundary CFT and W_{AB} is simply connected. Nevertheless, a thread configuration that maximizes the number of threads intersecting γ_{AB} and Γ can always be constructed according to the first prescription that we gave above. This is because the inequalities $S(AB) \leq S(A) + S(B)$, $I(A : B)/2 \leq E_W(A : B) \leq \min\{S(A), S(B)\}$, and $S(AB) \geq |S(A) - S(B)|$ guarantee that the right number of thread fragments can always be connected across γ_{AB} or Γ and attached to A and B as we described to construct the maximizing configuration. A qualitatively different example is the case where W_{AB} is disconnected, i.e., γ_{AB} is just the union of γ_A and γ_B . In this case $E_W(A : B) = 0$ and so there is no surface Γ to saturate. Another example is when AB is an entire CFT boundary but ρ_{AB} is a mixed state, such as the case of a single-sided mixed state black hole shown in figure 5. In this latter example, threads that leave W_{AB} terminate on the black hole horizon.⁶

Regulating lengths of surfaces near the boundary using a cutoff ϵ at finite coordinate distance into the bulk, the length of a generic boundary region in Planck units scales as ϵ^{-1} , while boundary- and bulk-anchored surfaces scale as $\log \epsilon$ and ϵ^0 , respectively. Thus, there will always be a great excess of available bit threads anchored to A and B to saturate the flow through the RT and entanglement wedge surfaces. Viewing the bit thread configurations as integral curves specified by a choice of vector field in the bulk, we must guarantee that the fields are chosen orthogonal to both the entanglement wedge

⁶Alternatively, we can think of the threads as crossing through a wormhole and terminating in another asymptotic boundary in which the black hole is purified.

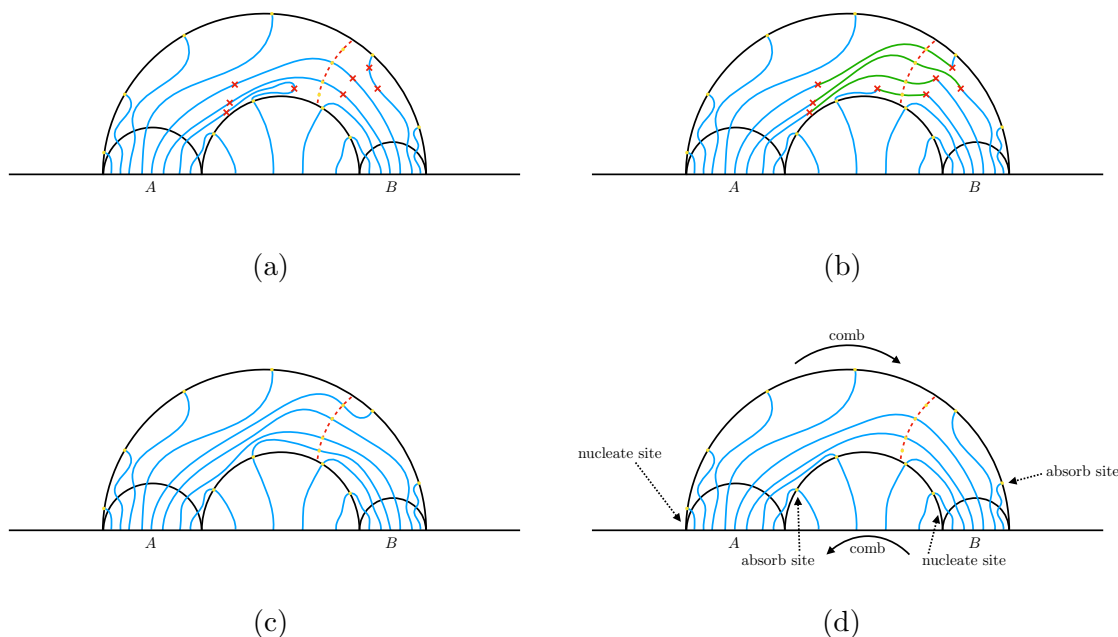


Figure 4. Starting with a thread configuration that saturates the number of threads intersecting γ_{AB} , γ_A , and γ_B , as shown in (a) and described in ref. [14], the goal is to construct a configuration like (c), which also saturates Γ . (The parts of threads that cross the exterior component of γ_{AB} and that lie outside the entanglement wedge have been suppressed in these diagrams.) Diagram (b) illustrates how this can be done by cutting and gluing threads. First, with a UV cutoff in place, it is helpful to think of there being a finite number of sites (yellow dots) on both γ_{AB} and Γ that threads must intersect. Initially, only $I(A : B)/2$ sites on Γ are filled by threads, which form a tube running through Γ . Divide the remaining $E_W(A : B) - I(A : B)/2$ sites into a group of n_{hi} sites above the tube and n_{lo} sites below the tube. (In two spatial dimensions the division is unique, but in higher dimensions there may be freedom in this choice.) Then, cut the $I(A : B)/2$ threads that cross Γ , as well as n_{hi} threads that intersect the exterior component of γ_{AB} on the side adjacent to B and n_{lo} threads that intersect the interior component of γ_{AB} on the side adjacent to A . The locations of the cuts are indicated by red crosses in (a). Finally, glue adjacent cut threads together as shown in (b) to arrive at the maximizing configuration (c). This cutting-and-gluing procedure is equivalent to “combing” the original thread configuration, as depicted in (d). Combing means dragging n_{hi} threads that intersect sites on the A -adjacent side of the exterior component of γ_{AB} over to sites on the B -adjacent side, and vice-versa for n_{lo} sites on the interior component of γ_{AB} . The only additional subtlety is that, with a UV cutoff in place, we must think of n_{hi} threads being nucleated from the UV where the exterior component of γ_{AB} meets A and n_{hi} sites being absorbed by the UV where the exterior component of γ_{AB} meets B . The same is true for n_{lo} threads hitting the interior component, with A and B flipped. Although the diagrams as we have drawn them here are directly reflective of $\text{AdS}_3/\text{CFT}_2$, we do not believe that there are any barriers to combing and cutting-and-gluing in arbitrary dimensions.

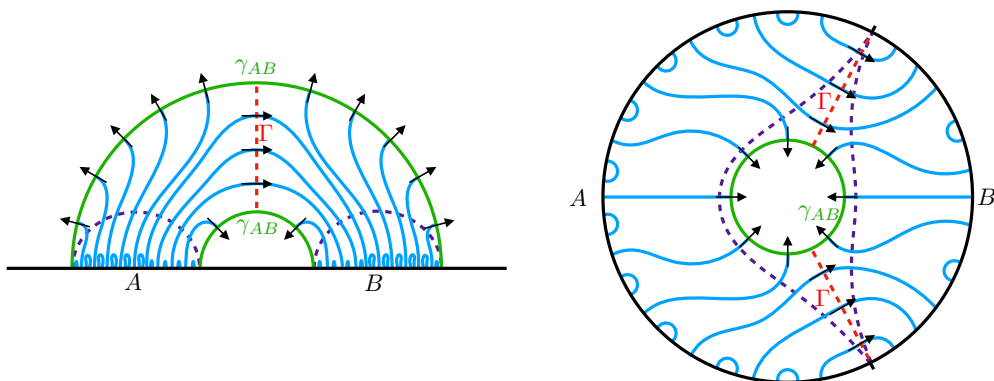


Figure 5. Representative examples of bit thread configurations simultaneously maximizing the flow through the entanglement wedge surface Γ (red dashed line) and RT surface γ_{AB} (green line). The vector fields specifying the flow are indicated on γ_{AB} and Γ by arrows, and the individual RT surfaces for A and B are depicted with purple dashed lines. Left: A and B are two surfaces defining a proper subregion of a single boundary, on which the CFT state is pure. Right: A and B partition an entire boundary, on which is defined a mixed CFT state, resulting in a horizon in the bulk.

cross section Γ and the RT surface γ_{AB} . As mentioned above, this is always possible except at $\Gamma \cap \gamma_{AB}$, which introduces errors only of measure-zero in the net flux. See figure 5 for representative examples.

The possible types of bit threads in a γ_{AB} - and Γ -saturating configuration are shown in figure 6. Here, we have only indicated threads that intersect W_{AB} . Although for a general thread configuration like that of figure 2 we could have threads that cross the RT surface multiple times, we can without loss of generality take such threads to be absent in a configuration (like that of figure 6) that obeys the maximization condition. This can be understood as follows. In the maximum-flux configuration, $|\gamma_{AB}|/4G_N\hbar$ threads cross γ_{AB} and are anchored to AB . Bit threads that pass through γ_{AB} an even number of times contribute zero net flow out of γ_{AB} , while those that pass through an odd number of times contribute the same net flow as those that pass through once. Hence, we can take the threads that pass through γ_{AB} to do so exactly once, without decreasing the maximum net flow out of AB . By the same token, since we are considering bit threads that simultaneously saturate the entanglement wedge cross section Γ , we can take the bit threads to cross Γ at most once, and there are $E_W(A : B)$ bit threads that do so.

Moreover, threads that start on A and end on A (or start on B and end on B) are not especially meaningful for our analysis. According to the maximization condition, there must be $|\gamma_{AB}|/4G_N\hbar$ threads passing through γ_{AB} and $E_W(A : B)$ threads passing through Γ . However, since $|AB|$ can be large compared to E_W and $|\gamma_{AB}|$ (but finite, once a UV cutoff is in place), it is possible to include some additional threads that start and terminate on A or B , while still retaining our max-flow conditions. However, such threads would necessarily cross either γ_{AB} or Γ an even number of times (including zero), and hence would contribute nothing to the thread configuration that we are maximizing. That is, such threads will not be relevant for any of the mutual informations we consider in our later calculations, so we suppress such extraneous threads from all subsequent figures and discussion.




Figure 6. The types of bit threads that can cross the RT surface of AB for the configuration that maximizes the flow out of AB as well as the flux through the entanglement wedge cross section Γ (red dashed line). We have highlighted threads that have one end in A or B and cross both Γ and the RT surface in green or blue, respectively.

Thus, we can take all bit threads that have some segment within the entanglement wedge to have an end in A , an end in B , or both. We will not depict possible threads that do not enter the entanglement wedge at all; according to the above discussion these should be associated with the subfactor $\mathcal{H}_{(AB)^c}/\mathcal{H}_{(AB)'}$.

3.3 Factorization and derivation of the upper bound

As reviewed above, $E_W(A : B)$ can be computed simply by counting the number of threads, in a Γ -saturating configuration, that cross the entanglement wedge cross section. Such threads can be categorized by whether they remain within the entanglement wedge connecting A and B , or instead start in A and leave through the RT surface γ_{AB} , or finally terminate in B having come through γ_{AB} . If, as suggested above, we identify threads that leave γ_{AB} as tags of the degrees of freedom in $(AB)'$, which purify ρ_{AB} , then the counting of threads that leave γ_{AB} computes various conditional mutual informations among A , B , and $(AB)'$.

Conditional mutual information (CMI) is a three-party entropic measure, defined as

$$I(A : B|C) = S(AC) + S(BC) - S(C) - S(ABC). \quad (3.3)$$

Similarly to how the number of orange-type threads connecting A and B in figure 6 gives $I(A : B)/2$, more generally, counting threads computes CMIs [23]. For example,

$$\begin{aligned} \frac{1}{2}I(A : (AB)'|B) &= \frac{1}{2} [S(AB) + S((AB)'B) - S(B) - S(AB(AB)')] \\ &= \frac{1}{2} [S(AB) + S(A) - S(B)] \\ &= S(A) - \frac{1}{2}I(A : B) \end{aligned}$$

$$\begin{aligned}
 &= (\# \text{ magenta, orange, green threads}) - (\# \text{ orange threads}) \\
 &= (\# \text{ magenta, green threads}).
 \end{aligned}$$

In the manipulations above, we used the fact that the geometric purification $|\Psi\rangle_{AB(AB)'}$ is pure⁷ to go from the first line to the second line.

Our crucial assumption will be to now postulate a factorization of $(AB)'$ into A' and B' according to whether threads that intersect γ_{AB} also intersect Γ and whether they are anchored to A or B . In other words, we suppose that $(AB)'$ factorizes into A' and B' such that, referring to figure 6, the CMIs among A , B , A' , and B' are given by the following thread counts:

- The total number of threads (crossing Γ) that pass from A to B counts $I(A : B)/2$.
- The number of green threads, which start in A , cross Γ , and leave the RT surface, counts $I(A : B'|B)/2$.
- The number of blue threads, which start in B , cross Γ , and then leave the RT surface, counts $I(B : A'|A)/2$.
- The number of magenta threads, which start in A and then leave the RT surface without passing through Γ , counts $I(A : A'|B)/2$.
- The number of purple threads, which start in B and leave the RT surface without passing through Γ , counts $I(B : B'|A)/2$.

We take these identifications to *define* A' and B' . A priori it is not clear that, for any choice of state ρ_{AB} , there always exists a pure state $|\Psi\rangle_{ABA'B'}$ whose CMIs are in fact equal to the numbers of colored threads as described above in a given bit thread configuration. In the broader case, with $I(A : B'|B)/2 = c_1$ and $I(B : A'|A)/2 = c_2$ for arbitrary c_1, c_2 , such pure states probably do not exist in general. In this section we *assume* that when ρ_{AB} is dual to an entanglement wedge and the purification is geometric, such a factorization of $(AB)'$ into A' and B' always exists.

Another way of phrasing our assumption is as follows. Let \mathcal{H} be the Hilbert space of the full CFT and $|\Psi\rangle \in \mathcal{H}$ be geometric. For any boundary subregion R , we can factorize $\mathcal{H} = \mathcal{H}_R \otimes \mathcal{H}_{R^c}$. There is always a state-dependent factorization $\mathcal{H}_{R^c} \cong \mathcal{H}_{R'} \otimes \mathcal{H}_Y$ such that $|\Psi\rangle = |\psi\rangle_{RR'} \otimes |\zeta\rangle_Y$. In this example, $|\psi\rangle_{RR'}$ corresponds to $|\Psi\rangle_{AB(AB)'}$ above. The state dependent factorization in question is precisely that of a partial entanglement distillation, which yields a factorization of exactly this form. The nontrivial assumption is that $(AB)'$ further factorizes to match the CMIs as counted by bit threads.

Specifically, we can consider $\mathcal{H}_{A'B'}$ as a subspace of $\mathcal{H}_{(AB)^c}$ because the bit threads pick out a subset of $(AB)^c$ that is entangled with AB , so there should exist a subspace of $(AB)^c$ that is unentangled with AB . This subspace will not be geometric, as any geometric

⁷Earlier, we defined the geometric purification $|\Psi\rangle_{AB(AB)^c}$ on one or more full boundaries with $\mathcal{H}_{(AB)^c} \cong \mathcal{H}_{(AB)'} \otimes \mathcal{H}_Y$, but because Y has no entanglement with AB , $\rho_{AB(AB)'}$ is pure and hence can be associated with a state vector.

subset of $(AB)^c$ will have some entanglement with AB , but there is no requirement in our construction that $(AB)'$ be a geometric boundary subregion, only that its support be contained within $(AB)^c$. Therefore, there exists a unitary transformation acting on $(AB)^c$ that takes the state $\rho_{(AB)^c}$ to $\rho_{A'B'} \otimes |\Phi\rangle\langle\Phi|$, where $|\Phi\rangle$ is a state in a subspace of $\mathcal{H}_{(AB)^c}$ representing the degrees of freedom that are unentangled with AB .

The best way to visualize this is as a partial entanglement distillation or an entanglement distillation with a partial inversion. One can first distill the entanglement between AB and $(AB)^c$; this will yield a product state of Bell pairs between a subset of AB and $A'B'$, and two unentangled pure states. Now, one can simply invert the half of the entanglement distillation on the AB portion to recover ρ_{AB} . As in appendix A, it is necessary for the entanglement cost of ρ_{AB} to equal the distillable entanglement of ρ_{AB} , as is guaranteed in ref. [33]. This form of partial entanglement distillation is made possible holographically primarily because of the holographic sandwiching result of ref. [34]. This (clearly state-dependent) procedure will then yield the state described in the previous paragraph.

The existence of a bit thread configuration might be interpreted as providing evidence that the state-dependent sub-factorization $(AB)' = A'B'$ can be achieved in the holographic case, but it would be nice to have a more concrete construction of such systems. We note that the existence of a surface-state correspondence [35] would immediately imply that such systems exist: in this case we could directly identify local purifying degrees of freedom living on the RT surface. We discuss this point further in section 5 below.

With these caveats, since E_W is given by the number of threads that cross the entanglement wedge cross section, our assumption then implies that

$$E_W(A : B) = \frac{1}{2} [I(A : B) + I(A : B'|B) + I(B : A'|A)]. \quad (3.4)$$

It is also true for general pure states $|\Psi\rangle_{ABA'B'}$ that

$$I(A : B) + I(A : B'|B) + I(B : A'|A) = 2S(AA') - I(A' : B'). \quad (3.5)$$

However, just as $I(A : B)/2$ is given by the number of threads connecting A to B in figure 6, $I(A' : B')/2$ must be identified with threads passing from one portion of $(AB)^c$ to another. But we argued above that these threads do not contribute to the purification of ρ_{AB} ; they live in the Y factor of the purifying geometry, i.e., in $\mathcal{H}_{(AB)^c}/\mathcal{H}_{A'B'}$. Hence, without loss of generality, we may take any bit thread configuration that saturates the RT surface and remove any bit threads that do not enter W_{AB} .

Therefore, it is true here that

$$I(A : B) + I(A : B'|B) + I(B : A'|A) = 2S(AA'). \quad (3.6)$$

Inserting eq. (3.4), we immediately have that

$$E_W(A : B) = S(AA') \quad (3.7)$$

for the choice of A' defined by the bit thread configuration we have specified. By the definition in section 2, $E_P(A : B) \leq S(AA')$ for any particular purification, so we have

$$E_P(A : B) \leq E_W(A : B). \quad (3.8)$$

4 Lower bound on E_P

Let us take our entanglement wedge for AB and write a purification as $|\psi\rangle_{AA'BB'Y}$, where we have padded with enough unentangled ancillae that the dimension of $|\psi\rangle$ has the dimension of a full boundary CFT with some cutoff. In particular, let us choose $|\psi\rangle$ to be a state that achieves the infimum of $S(AA')$. (Once a cutoff is in place, which makes the boundary theory finite-dimensional (cf. section 2.1), the infimum may always be achieved.) Now, the set of CFT states $|\chi_i\rangle$ on the entire boundary that are dual to a classical holographic bulk form an overcomplete basis for all CFT states [36]. By a slight generalization, we should be able to express $|\psi\rangle$ as some superposition

$$|\psi\rangle = \sum_{i=1}^M \alpha_i |\phi_i\rangle, \tag{4.1}$$

where the $|\phi_i\rangle$ are a subset of the $|\chi_i\rangle$ for which the reduced density matrix on AB is fixed, $\text{Tr}_{(AB)^c} |\phi_i\rangle\langle\phi_i| = \rho_{AB}$ for all i . That is, we have written the purification as a superposition over entire classical bulk geometries dual to $|\phi_i\rangle$, such that each has the same geometry in the entanglement wedge.

We can argue for the existence of the decomposition (4.1) as follows. Starting from some particular geometric purification of ρ_{AB} , we are free to add pairs of black holes connected by wormholes⁸ to the bulk geometry outside of W_{AB} . As this is a local bulk modification on only the purifying subsystem (up to the caveats discussed in footnote 4) this action modifies the boundary state without changing the fact that it is a purification of ρ_{AB} . (In the case where AB comprises an entire CFT boundary, with thermal density matrix dual to a black hole in the bulk, this construction simply amounts to gluing on different multiboundary wormholes behind the horizon.) Let us call the dual boundary states formed by this procedure $|\xi_j\rangle \subset \{|\phi_j\rangle\}$, i.e., these geometries are a subset of those given by the set of all geometric purifications of ρ_{AB} .⁹ Importantly, although general purifications need not all live in the same Hilbert space, the particular dual states $|\xi_j\rangle$ all live in a single Hilbert space, that of a large- N holographic CFT, which also contains the geometric purification defined in the previous subsection. (In the case where ρ_{AB} describes

⁸Altering an asymptotically-AdS geometry by producing an entangled pair of black holes changes the topology of the spacetime. This is certainly an allowed process in quantum gravity [37–40]; for example, we can consider inserting initial excitations that collide to produce a pair of entangled black holes, i.e., the time-reversal of the process by which a pair of entangled black holes evaporate. Hence asymptotically-AdS spacetimes in which this topology-changing process takes place should still have a good dual description in the large- N CFT; they can be produced simply by specifying appropriate initial conditions, with no need to deform the theory. However, the topology change is a fundamentally nonperturbative process from the point of view of a classical bulk description as a curved-space QFT, and thus we should not expect this new spacetime to be in the same code subspace [25] as one in which the process never occurs.

⁹As noted in footnote 3, this black hole gluing construction is straightforward in asymptotically-AdS₃ geometries, due to the vanishing of the Weyl tensor. More generally, the construction of Engelhardt and Wall [31] (see also ref. [19]) allows us to glue black hole handles onto geometries in higher dimensions; though such states may not be dynamically stable [41], this will not matter for our purposes here, since we merely require the existence of such a geometry as a locally time-symmetric solution of Einstein’s equations on a Cauchy slice.

a thermal state, the Hilbert space is instead that of two CFTs, and includes the appropriate thermofield double states.)

In particular, we are free to choose the masses and entanglement structure of such black holes, so long as we do not change the purity of the overall state. These black holes should have support over any complete basis of $\mathcal{H}_{(AB)^c}$, by the same reasoning as in the case of thermofield double states of differing temperatures [42, 43], as we can include all black holes from those of minimal size to ones that fill up a large portion of the purifying bulk region. Concretely, considering a basis $|\omega_k\rangle$ for $\mathcal{H}_{(AB)^c}$ and writing $\text{Tr}_{AB} |\xi_j\rangle\langle\xi_j| = \sigma_j$, for all j we should be able to write

$$\sigma_j = \sum_k c_{jk} |\omega_k\rangle\langle\omega_k| + \sum_{k \neq k'} c_{jkk'} |\omega_k\rangle\langle\omega_{k'}|, \tag{4.2}$$

where the second sum containing the off-diagonal terms is exponentially suppressed relative to the first sum over the diagonal terms,¹⁰ and c_{jk} is (to good approximation) invertible.¹¹ We expect distinct σ_j , which correspond to distinct classical metrics outside of W_{AB} , to be orthogonal up to exponential corrections, i.e., $\text{Tr} \sigma_j \sigma_{j'} \simeq \delta_{jj'}$ (by the same reasoning that implies that distinct thermofield double states have exponentially small overlap [45]). Hence, a superposition of the $|\xi_j\rangle$ will retain the same reduced density matrix ρ_{AB} on AB . The linear independence of the c_{jk} weighting allows the $|\omega_k\rangle\langle\omega_k|$ themselves to be isolated by performing Gauss-Jordan elimination over the black hole geometries: $|\omega_k\rangle\langle\omega_k| \simeq \sum_j a_{jk} \sigma_j$, where $a = c^{-1}$. Hence, the number of linearly independent black hole states among the $|\xi_j\rangle$ is at least $\dim \mathcal{H}_{(AB)^c}$, so for some α_j we have

$$|\psi\rangle = \sum_j \alpha_j |\xi_j\rangle, \tag{4.3}$$

and therefore the $|\xi_j\rangle$ are a valid choice of $|\phi_i\rangle$ in the sum in eq. (4.1).

It is worth emphasizing at this point that this argument suggests a subregion version of the fact that geometric states form an overcomplete basis of the entire Hilbert space. This idea is an intuitive generalization, but we have motivated it above as we are unaware of this statement's appearance in extant literature.

Now we wish to compute $S(AA')$ in $|\psi\rangle$ and determine its minimum value. Let us write $S_i = \text{Tr}_{(AA')^c} |\phi_i\rangle\langle\phi_i|$, so S_i is the entanglement entropy of AA' in the geometry i . In each individual geometry, we can write down a concrete bit thread configuration, where each bit thread is boundary-anchored, so we can compute S_i by counting bit threads. As shown in ref. [45], we may write $S(AA')$ as a sum $\sum_i |\alpha_i|^2 S_i$ under certain assumptions, which we will discuss below. Thus, we can find a lower bound on $S(AA')$ by simply lower-bounding S_i . We can lower-bound S_i using a bit thread configuration in an everywhere-classical

¹⁰We expect this block-diagonal assumption to be true at leading order in N by the eigenstate thermalization hypothesis [44].

¹¹Note in particular that we are not claiming that the $|\xi_j\rangle$ form a complete basis for the entire CFT Hilbert space; in particular, we will not be able to reach states that have reduced density matrices on AB orthogonal to ρ_{AB} .

geometry. Putting this all together, we have:

$$\begin{aligned}
 E_P(A : B) &= S(\text{Tr}_{BB'Y} |\psi\rangle\langle\psi|) \\
 &= \sum_{i=1}^M |\alpha_i|^2 S_i + S_{\text{mix}} + \text{subleading} \\
 &\geq \min_i S_i + S_{\text{mix}} + \text{subleading} \tag{4.4} \\
 &= E_W(A : B) + \min_i \frac{1}{2} I(A' : B')_i + S_{\text{mix}} + \text{subleading} \quad (\text{from threads}) \\
 &\geq E_W(A : B),
 \end{aligned}$$

where S_{mix} is the entropy of mixing, $-\sum_i |\alpha_i|^2 \log |\alpha_i|^2$. An important caveat is that the superposition formula that we discuss above applies only when the number of terms M in the superposition is small relative to $e^{O(c)}$. The necessity of this restriction can be understood from the fact that any product state can be written as a superposition of a large number of entangled geometric states, exploiting the overcompleteness of the basis, as in, e.g., ref. [46]. Therefore, an assumption must be made here to conclude the argument above: the number of terms in the superposition must still be small relative to when the approximation breaks down. An argument supporting this assumption follows from our reasoning for the existence of the decomposition in eq. (4.1) itself, using the construction involving black holes and subsequent row reduction of the basis states. That construction suggests that the number of terms in the superposition (4.1) will scale at most as $\dim \mathcal{H}_{(AB)^c}$. Because the total set of states for the entire boundary CFT scales as $\dim \mathcal{H}_{AB(AB)^c} \sim e^{O(c)}$, there will be a suppression in the entropy-of-mixing term relative to the leading term that goes like

$$\frac{S_{\text{mix}}}{E_W(A : B)} \lesssim \frac{\log M}{c} \lesssim \frac{\log \dim \mathcal{H}_{(AB)^c}}{\log \dim \mathcal{H}_{AB(AB)^c}} < 1, \tag{4.5}$$

where the first inequality comes from $E_W \sim c$ [47] and the upper bound on the Shannon entropy, $S_{\text{mix}} \leq \log M$, the second inequality comes from the definition of the Hilbert space on $(AB)^c$ for CFTs, and the final inequality is definitionally true. When the number of terms M is small compared to $\dim \mathcal{H}_{(AB)^c}$, i.e., when $|\psi\rangle$ can be formed by the superposition of a small number of classical states, then the second inequality in eq. (4.5) can be strongly satisfied. In eq. (4.5), we are working in the formalism of a UV-regulated CFT, so that all Hilbert space dimensions are finite, making all the ratios well defined; holographically, this just corresponds to imposing a cutoff at fixed radial coordinate near the boundary, which also makes all of the geometric quantities we consider finite.

One potentially illuminating way of thinking about this is to consider the reason why the block-diagonal approximation of $|\psi\rangle$ is appropriate for sufficiently small superpositions. The reason is because the off-diagonal terms are suppressed by e^c relative to the sum of the diagonal terms of $|\psi\rangle\langle\psi|$. This suppression only becomes weak enough to merit consideration of the off-diagonal terms when the number of nontrivial diagonal terms approaches e^c ; notably, if the number is only e^{kc} for $k < 1$, there should still be a relative exponential

suppression of the off-diagonal elements by $e^{(k-1)c}$. This suggests that the inequality in eq. (4.5) may in fact be even stronger than it appears, as it deals in terms of entropies as opposed to Hilbert space dimensionality.

Putting together eqs. (3.8) and (4.4), we obtain

$$E_P(A : B) = E_W(A : B). \tag{4.6}$$

We have thus proven the $E_P = E_W$ conjecture, under the assumption that we can associate geometric purifications with bit thread configurations in the sense of section 3 and, as motivated in this section, that our state $|\psi\rangle$ realizing the optimal purification can be written as a superposition of fewer than $\dim \mathcal{H}$ geometric states, each of which has fixed reduced density matrix on AB corresponding to the entanglement wedge geometry.

5 Discussion

It should be made clear that some of the assumptions in section 4 have particularly strong support in three-dimensional gravity, where both a) the triviality of the Weyl tensor straightforwardly allows for the gluing of arbitrary black hole handles and b) the validity of Cardy’s formula for density of states [42] is most compelling. We nevertheless expect some, probably more complicated, version of these arguments to persist in higher dimensions;¹² we leave such investigations to future work. It would be a highly surprising development if subregion overcompleteness of geometric states were true only in three dimensions.

It is also worth stressing that the program advocated in this paper is not equivalent to the surface/state correspondence [35]. Nowhere in the present work was it necessary to localize Hilbert space subfactors to the RT surface or indeed to any bulk surface; the entire argument was made from the perspective of boundary Hilbert space subfactors. However, if we do assume the surface/state correspondence, then we do not need to extend the bit threads to an entire bulk geometry, since we can then localize the purifying degrees of freedom to the boundary of the entanglement wedge. In this case, the identification of the purification becomes trivial, with A' (or B') corresponding to the portions of γ_{AB} between A (or B , respectively) and Γ . The connection between the surface/state correspondence and the $E_P = E_W$ conjecture was first pointed out in ref. [12].

We also expect that the bit threads justification for $E_P = E_W$ will extend to the multipartite generalizations of entanglement of purification studied in refs. [15–18]. In particular, because we are classifying flows through surfaces, the same analysis should follow in these cases. In the multipartite generalization, the property that the minimal polytope in the bulk is inscribed within the RT surfaces has the consequence that, in the bit thread construction, all bit threads must cross the minimal polytope an even number of times (including zero), which allows for the RT surfaces to be completely partitioned into the A'_i . We leave investigation of the bit thread picture for multipartite entanglement of purification to future work.

¹²Indeed, as we have noted above, the gluing construction can be generalized using the formalism of refs. [19, 31].

Acknowledgments

We thank Charles Cao, Illan Halpern, Matt Headrick, Yasunori Nomura, Nico Salzetta, and Mark Van Raamsdonk for useful discussions and comments. We are especially grateful to Jamie Sully for numerous fruitful discussions and initial collaboration. N.B. is supported by the National Science Foundation under grant number 82248-13067-44-PHPXH, by the Department of Energy under grant number DE-SC0019380, and by New York State Urban Development Corporation Empire State Development contract no. AA289. A.C.-D. is supported by the KU Leuven C1 grant ZKD1118 C16/16/005, the National Science Foundation of Belgium (FWO) grant G.001.12 Odysseus, and by the European Research Council grant no. ERC-2013-CoG 616732 HoloQosmos. J.P. is supported in part by the Simons Foundation and in part by the Natural Sciences and Engineering Research Council of Canada. G.N.R. is supported by the Miller Institute for Basic Research in Science at the University of California, Berkeley.

A Holographic purification of minimal dimension

Consider a purification $|\psi\rangle_{AA'BB'}$ of a state ρ_{AB} — in a holographic theory with a UV cutoff— that realizes the infimum of $S(AA')$ over all purifications of ρ_{AB} . In this appendix we argue that $|\psi\rangle_{AA'BB'}$ can be compressed in the dimensionality of $\mathcal{H}_{A'} \otimes \mathcal{H}_{B'}$ to a new purification $|\phi\rangle_{AA'BB'}$ such that $\dim(\mathcal{H}_{A'} \otimes \mathcal{H}_{B'}) = e^{S(\rho_{AB})}$, the minimal possible dimension of a purification of ρ_{AB} .

We begin by noting that a distillation of Bell pairs between AA' and BB' for $|\psi\rangle_{AA'BB'}$ is an LOCC procedure, so it preserves $S(AA')$. If we view A , A' , B , and B' as each being some large tensor product over a number of qubits, where we imagine that each qubit is either in a Bell state or unentangled (i.e., ignoring issues of multipartite entanglement), then there should be no pairs of qubits, one in A' and one in B' , that are in a Bell state. In particular, this is true because, were such pairs of qubits to exist, one could decrease $S(AA')$ by simply excising such pairs of qubits from $|\psi\rangle_{AA'BB'}$, in contradiction with the hypothesis that $|\psi\rangle_{AA'BB'}$ realizes the infimum of $S(AA')$. That is, we should have $I(A' : B') = 0$. More generally, if there were multipartite entanglement among subfactors of A , A' , B , and B' , the purifying system could be re-factored to eliminate any contribution to $I(A' : B')$. For illustration, supposing that there were tripartite entanglement between some subfactors $a \subset A$, $a' \subset A'$, and $b' \subset B'$, then redefining A' to include a' and b' eliminates the possible contribution to $I(A' : B')$, while also lowering $S(AA')$. Again, if this were possible the initial partition in fact could not have been optimal.

We now perform the distillation procedure, which results in a tensor product of three pure states:

$$|\psi\rangle_{AA'BB'} \in \mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_{A'} \otimes \mathcal{H}_{B'} \xrightarrow{\text{LOCC}} |\phi\rangle_{AA'BB'} = |\phi\rangle_{A_0A'_0} \otimes |\phi\rangle_{B_0B'_0} \otimes |\phi\rangle_{\text{bell}}, \quad (\text{A.1})$$

where

$$\begin{aligned}
 \mathcal{H}_A &= \mathcal{H}_{A_0} \otimes \mathcal{H}_{A_{\text{bell}}} \\
 \mathcal{H}_B &= \mathcal{H}_{B_0} \otimes \mathcal{H}_{B_{\text{bell}}} \\
 \mathcal{H}_{A'} &= \mathcal{H}_{A'_0} \otimes \mathcal{H}_{A'_{\text{bell}}} \\
 \mathcal{H}_{B'} &= \mathcal{H}_{B'_0} \otimes \mathcal{H}_{B'_{\text{bell}}}
 \end{aligned}
 \tag{A.2}$$

and where

$$\begin{aligned}
 |\phi\rangle_{\text{bell}} &\in \mathcal{H}_{A_{\text{bell}}} \otimes \mathcal{H}_{B_{\text{bell}}} \otimes \mathcal{H}_{A'_{\text{bell}}} \otimes \mathcal{H}_{B'_{\text{bell}}} \\
 |\phi\rangle_{A_0 A'_0} &\in \mathcal{H}_{A_0} \otimes \mathcal{H}_{A'_0} \\
 |\phi\rangle_{B_0 B'_0} &\in \mathcal{H}_{B_0} \otimes \mathcal{H}_{B'_0}.
 \end{aligned}
 \tag{A.3}$$

Note that, at this point, none of the Hilbert space factors are “prunable,” in the sense that none of the factors of $|\phi\rangle_{AA'BB'}$ lives exclusively in $\mathcal{H}_{A'_0}$ or $\mathcal{H}_{B'_0}$ (and we recall that we are not allowed to delete any data about ρ_{AB} , lest we jeopardize the recoverability of our original density matrix).

Next, let us distill twice more on $|\phi\rangle_{A_0 A'_0}$ and $|\phi\rangle_{B_0 B'_0}$:

$$\begin{aligned}
 |\phi\rangle_{A_0 A'_0} &\xrightarrow{\text{LOCC}} |\phi\rangle_{A_{00}} \otimes |\phi\rangle_{A'_{00}} \otimes |\phi\rangle_{A_0 A'_0, \text{bell}} \\
 |\phi\rangle_{B_0 B'_0} &\xrightarrow{\text{LOCC}} |\phi\rangle_{B_{00}} \otimes |\phi\rangle_{B'_{00}} \otimes |\phi\rangle_{B_0 B'_0, \text{bell}},
 \end{aligned}
 \tag{A.4}$$

where

$$\begin{aligned}
 \mathcal{H}_{A_0} &= \mathcal{H}_{A_{00}} \otimes \mathcal{H}_{A_0, \text{bell}} \\
 \mathcal{H}_{B_0} &= \mathcal{H}_{B_{00}} \otimes \mathcal{H}_{B_0, \text{bell}} \\
 \mathcal{H}_{A'_0} &= \mathcal{H}_{A'_{00}} \otimes \mathcal{H}_{A'_0, \text{bell}} \\
 \mathcal{H}_{B'_0} &= \mathcal{H}_{B'_{00}} \otimes \mathcal{H}_{B'_0, \text{bell}}
 \end{aligned}
 \tag{A.5}$$

and where

$$\begin{aligned}
 |\phi\rangle_{A_0 A'_0, \text{bell}} &\in \mathcal{H}_{A_0, \text{bell}} \otimes \mathcal{H}_{A'_0, \text{bell}} \\
 |\phi\rangle_{B_0 B'_0, \text{bell}} &\in \mathcal{H}_{B_0, \text{bell}} \otimes \mathcal{H}_{B'_0, \text{bell}} \\
 |\phi\rangle_{A_{00}} &\in \mathcal{H}_{A_{00}} \\
 |\phi\rangle_{B_{00}} &\in \mathcal{H}_{B_{00}} \\
 |\phi\rangle_{A'_{00}} &\in \mathcal{H}_{A'_{00}} \\
 |\phi\rangle_{B'_{00}} &\in \mathcal{H}_{B'_{00}}.
 \end{aligned}
 \tag{A.6}$$

Note that these distillations will also not affect $S(AA')$, since the LOCCs in eq. (A.4) are acting only on subfactors of the Hilbert space that are unused in calculating $S(AA')$, as the latter depends only on $|\phi\rangle_{\text{bell}}$. See figure 7 for an illustration of the Hilbert space decomposition and entanglement distillation in eqs. (A.1) through (A.6).

Now, we are in a position to “prune” the Hilbert space factors $\mathcal{H}_{A'_{00}}$ and $\mathcal{H}_{B'_{00}}$, since these have been rendered superfluous. Let us then enumerate the total remaining Hilbert space dimensions of the “primed” factors. We note that the remaining “primed” factors of the Hilbert space now each contain a state that is maximally mixed and purified by some analogous unprimed factor. This means that the Hilbert space dimension of the remaining




Figure 7. Hilbert space decomposition and entanglement distillation of AB and its purification.

primed factors (i.e., less $\mathcal{H}_{A'_{00}} \otimes \mathcal{H}_{B'_{00}}$) equals $e^{S(\rho_{AB})}$, because the state on the remaining primed factors is simply a maximally-mixed purification of a state σ_{AB} obtained by tracing out all the remaining primed factors from our state we have after all the LOCC steps, for which $S(\sigma_{AB}) = S(\rho_{AB})$.

We now note several additional facts. First, holographically, ρ_{AB} is recoverable to leading order from σ_{AB} via a suitable unitary transformation, which follows from the fact that $S(\sigma_{AB}) = S(\rho_{AB})$ to leading order by the Hayden-Swingle-Walter sandwiching result [34] (for a proof sketch, see section 2.3 of ref. [33]). It was also necessary that there existed no entanglement between the reduced states on $\mathcal{H}_{A'_{\text{bell}}}$ and $\mathcal{H}_{B'_{\text{bell}}}$, in order to guarantee the maximal mixing condition that led to our conclusion of minimality of dimension. Having started with a state $|\psi\rangle_{AA'BB'}$ with no mutual information between A' and B' (which we argued was a consequence of $|\psi\rangle_{AA'BB'}$ achieving the infimum of $S(AA')$), the Bell pair distillation gains no advantage by introducing correlation between A'_{bell} and B'_{bell} , so this condition should follow.

In addition, in order for ρ_{AB} to be recoverable after our LOCC steps, it is necessary for the entanglement cost of ρ_{AB} to equal the distillable entanglement of ρ_{AB} . Luckily, this equality was shown for holographic states (like ρ_{AB}) in ref. [33].

In summary, given a purification of the holographic state on AB , such that the purification minimizes $S(AA')$, we can exhibit a purification, with the same $S(AA')$, that also attains the minimal possible Hilbert space dimension.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] G. 't Hooft, *Dimensional reduction in quantum gravity*, *Conf. Proc. C* **930308** (1993) 284 [[gr-qc/9310026](https://arxiv.org/abs/gr-qc/9310026)] [[INSPIRE](https://inspirehep.net/literature/32133)].

- [2] L. Susskind, *The world as a hologram*, *J. Math. Phys.* **36** (1995) 6377 [[hep-th/9409089](#)] [[INSPIRE](#)].
- [3] R. Bousso, *The holographic principle*, *Rev. Mod. Phys.* **74** (2002) 825 [[hep-th/0203101](#)] [[INSPIRE](#)].
- [4] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Phys. Rev. Lett.* **96** (2006) 181602 [[hep-th/0603001](#)] [[INSPIRE](#)].
- [5] S. Ryu and T. Takayanagi, *Aspects of holographic entanglement entropy*, *JHEP* **08** (2006) 045 [[hep-th/0605073](#)] [[INSPIRE](#)].
- [6] A. Lewkowycz and J. Maldacena, *Generalized gravitational entropy*, *JHEP* **08** (2013) 090 [[arXiv:1304.4926](#)] [[INSPIRE](#)].
- [7] J.M. Maldacena, *The large N limit of superconformal field theories and supergravity*, *Int. J. Theor. Phys.* **38** (1999) 1113 [[hep-th/9711200](#)] [[INSPIRE](#)].
- [8] S.S. Gubser, I.R. Klebanov and A.M. Polyakov, *Gauge theory correlators from noncritical string theory*, *Phys. Lett. B* **428** (1998) 105 [[hep-th/9802109](#)] [[INSPIRE](#)].
- [9] E. Witten, *Anti-de Sitter space and holography*, *Adv. Theor. Math. Phys.* **2** (1998) 253 [[hep-th/9802150](#)] [[INSPIRE](#)].
- [10] O. Aharony, S.S. Gubser, J.M. Maldacena, H. Ooguri and Y. Oz, *Large N field theories, string theory and gravity*, *Phys. Rept.* **323** (2000) 183 [[hep-th/9905111](#)] [[INSPIRE](#)].
- [11] B.M. Terhal, M. Horodecki, D.W. Leung and D.P. DiVincenzo, *The entanglement of purification*, *J. Math. Phys.* **43** (2002) 4286 [[quant-ph/0202044](#)].
- [12] T. Takayanagi and K. Umemoto, *Entanglement of purification through holographic duality*, *Nature Phys.* **14** (2018) 573 [[arXiv:1708.09393](#)] [[INSPIRE](#)].
- [13] P. Nguyen, T. Devakul, M.G. Halbasch, M.P. Zaletel and B. Swingle, *Entanglement of purification: from spin chains to holography*, *JHEP* **01** (2018) 098 [[arXiv:1709.07424](#)] [[INSPIRE](#)].
- [14] C.A. Agón, J. De Boer and J.F. Pedraza, *Geometric aspects of holographic bit threads*, *JHEP* **05** (2019) 075 [[arXiv:1811.08879](#)] [[INSPIRE](#)].
- [15] N. Bao and I.F. Halpern, *Holographic inequalities and entanglement of purification*, *JHEP* **03** (2018) 006 [[arXiv:1710.07643](#)] [[INSPIRE](#)].
- [16] N. Bao and I.F. Halpern, *Conditional and multipartite entanglements of purification and holography*, *Phys. Rev. D* **99** (2019) 046010 [[arXiv:1805.00476](#)] [[INSPIRE](#)].
- [17] N. Bao, A. Chatwin-Davies and G.N. Remmen, *Entanglement of purification and multiboundary wormhole geometries*, *JHEP* **02** (2019) 110 [[arXiv:1811.01983](#)] [[INSPIRE](#)].
- [18] K. Umemoto and Y. Zhou, *Entanglement of purification for multipartite states and its holographic dual*, *JHEP* **10** (2018) 152 [[arXiv:1805.02625](#)] [[INSPIRE](#)].
- [19] S. Dutta and T. Faulkner, *A canonical purification for the entanglement wedge cross-section*, [arXiv:1905.00577](#) [[INSPIRE](#)].
- [20] J. Kudler-Flam, I. MacCormack and S. Ryu, *Holographic entanglement contour, bit threads and the entanglement tsunami*, *J. Phys. A* **52** (2019) 325401 [[arXiv:1902.04654](#)] [[INSPIRE](#)].
- [21] D.-H. Du, C.-B. Chen and F.-W. Shu, *Bit threads and holographic entanglement of purification*, [arXiv:1904.06871](#) [[INSPIRE](#)].

- [22] N. Jokela and A. Pönni, *Notes on entanglement wedge cross sections*, [arXiv:1904.09582](#) [[INSPIRE](#)].
- [23] M. Freedman and M. Headrick, *Bit threads and holographic entanglement*, *Commun. Math. Phys.* **352** (2017) 407 [[arXiv:1604.00354](#)] [[INSPIRE](#)].
- [24] S.X. Cui, P. Hayden, T. He, M. Headrick, B. Stoica and M. Walter, *Bit threads and holographic monogamy*, [arXiv:1808.05234](#) [[INSPIRE](#)].
- [25] D. Harlow, *The Ryu-Takayanagi formula from quantum error correction*, *Commun. Math. Phys.* **354** (2017) 865 [[arXiv:1607.03901](#)] [[INSPIRE](#)].
- [26] I. Bakhmatov, N.S. Deger, J. Gutowski, E. Ó Colgáin and H. Yavartanoo, *Calibrated entanglement entropy*, *JHEP* **07** (2017) 117 [[arXiv:1705.08319](#)] [[INSPIRE](#)].
- [27] R. Harvey and H.B. Lawson, *Calibrated geometries*, *Acta Math.* **148** (1982) 47.
- [28] H. Federer, *Real flat chains, cochains and variational problems*, *Indiana Univ. Math. J.* **24** (1974) 351.
- [29] G. Strang, *Maximal flow through a domain*, *Math. Prog.* **26** (1983) 123.
- [30] R. Nozawa, *Max-flow min-cut theorem in an anisotropic network*, *Osaka J. Math.* **27** (1990) 805.
- [31] N. Engelhardt and A.C. Wall, *Coarse graining holographic black holes*, *JHEP* **05** (2019) 160 [[arXiv:1806.01281](#)] [[INSPIRE](#)].
- [32] M. Fannes, *A continuity property of the entropy density for spin lattice systems*, *Commun. Math. Phys.* **31** (1973) 291.
- [33] N. Bao, G. Penington, J. Sorce and A.C. Wall, *Beyond toy models: distilling tensor networks in full AdS/CFT*, [arXiv:1812.01171](#) [[INSPIRE](#)].
- [34] P. Hayden, B. Swingle and M. Walter, forthcoming.
- [35] M. Miyaji and T. Takayanagi, *Surface/state correspondence as a generalized holography*, *PTEP* **2015** (2015) 073B03 [[arXiv:1503.03542](#)] [[INSPIRE](#)].
- [36] M. Botta-Cantcheff and P.J. Martínez, *Which quantum states are dual to classical spacetimes?*, [arXiv:1703.03483](#) [[INSPIRE](#)].
- [37] D. Garfinkle and A. Strominger, *Semiclassical Wheeler wormhole production*, *Phys. Lett.* **B 256** (1991) 146 [[INSPIRE](#)].
- [38] J. Maldacena and L. Susskind, *Cool horizons for entangled black holes*, *Fortsch. Phys.* **61** (2013) 781 [[arXiv:1306.0533](#)] [[INSPIRE](#)].
- [39] N. Bao, J. Pollack and G.N. Remmen, *Splitting spacetime and cloning qubits: linking no-go theorems across the ER=EPR duality*, *Fortsch. Phys.* **63** (2015) 705 [[arXiv:1506.08203](#)] [[INSPIRE](#)].
- [40] S.B. Giddings and R.A. Porto, *The gravitational S-matrix*, *Phys. Rev.* **D 81** (2010) 025002 [[arXiv:0908.0004](#)] [[INSPIRE](#)].
- [41] E. Witten and S.-T. Yau, *Connectedness of the boundary in the AdS/CFT correspondence*, *Adv. Theor. Math. Phys.* **3** (1999) 1635 [[hep-th/9910245](#)] [[INSPIRE](#)].
- [42] J.L. Cardy, *Operator content of two-dimensional conformally invariant theories*, *Nucl. Phys.* **B 270** (1986) 186 [[INSPIRE](#)].

- [43] J.M. Maldacena, *Eternal black holes in anti-de Sitter*, *JHEP* **04** (2003) 021 [[hep-th/0106112](#)] [[INSPIRE](#)].
- [44] M. Srednicki, *Chaos and quantum thermalization*, *Phys. Rev. E* **50** (1994) 888 [[cond-mat/9403051](#)].
- [45] A. Almheiri, X. Dong and B. Swingle, *Linearity of holographic entanglement entropy*, *JHEP* **02** (2017) 074 [[arXiv:1606.04537](#)] [[INSPIRE](#)].
- [46] M. Van Raamsdonk, *Building up spacetime with quantum entanglement*, *Gen. Rel. Grav.* **42** (2010) 2323 [[arXiv:1005.3035](#)] [[INSPIRE](#)].
- [47] J.D. Brown and M. Henneaux, *Central charges in the canonical realization of asymptotic symmetries: an example from three-dimensional gravity*, *Commun. Math. Phys.* **104** (1986) 207 [[INSPIRE](#)].