

Beyond toy models: distilling tensor networks in full AdS/CFT

Ning Bao,^a Geoffrey Penington,^b Jonathan Sorce^b and Aron C. Wall^c

^a*Berkeley Center for Theoretical Physics, University of California,
366 Le Conte Hall, Berkeley, CA 94720-7300, U.S.A.*

^b*Stanford Institute for Theoretical Physics, Stanford University,
382 Via Pueblo Mall, Stanford, CA 94305-4060, U.S.A.*

^c*Centre for Mathematical Sciences, Cambridge University,
Wilberforce Rd, Cambridge CB3 0WA, U.K.*

E-mail: ningbao75@gmail.com, geoffp@stanford.edu, jsorce@stanford.edu,
aroncwall@gmail.com

ABSTRACT: We present a general procedure for constructing tensor networks that accurately reproduce holographic states in conformal field theories (CFTs). Given a state in a large- N CFT with a static, semiclassical gravitational dual, we build a tensor network by an iterative series of approximations that eliminate redundant degrees of freedom and minimize the bond dimensions of the resulting network. We argue that the bond dimensions of the tensor network will match the areas of the corresponding bulk surfaces. For “tree” tensor networks (i.e., those that are constructed by discretizing spacetime with non-intersecting Ryu-Takayanagi surfaces), our arguments can be made rigorous using a version of one-shot entanglement distillation in the CFT. Using the known quantum error correcting properties of AdS/CFT, we show that bulk legs can be added to the tensor networks to create holographic quantum error correcting codes. These codes behave similarly to previous holographic tensor network toy models, but describe actual bulk excitations in continuum AdS/CFT.

By assuming some natural generalizations of the “holographic entanglement of purification” conjecture, we are able to construct tensor networks for more general bulk discretizations, leading to finer-grained networks that partition the information content of a Ryu-Takayanagi surface into tensor-factorized subregions. While the granularity of such a tensor network must be set larger than the string/Planck scales, we expect that it can be chosen to lie well below the AdS scale. However, we also prove a no-go theorem which shows that the bulk-to-boundary maps cannot all be isometries in a tensor network with intersecting Ryu-Takayanagi surfaces.

KEYWORDS: AdS-CFT Correspondence, Gauge-gravity correspondence, Conformal Field Theory

ARXIV EPRINT: [1812.01171](https://arxiv.org/abs/1812.01171)

Contents

1	Introduction	2
2	Tensor networks, entanglement distillation, and one-shot quantum information	4
2.1	Tensor networks	4
2.2	Entanglement distillation	7
2.3	Holographic one-shot entropies	9
3	Tree tensor networks from holographic entanglement distillation	12
3.1	Bipartite tensor networks	13
3.2	General tree networks	18
4	Bulk legs and holographic quantum error correction	26
5	Loop tensor networks from holographic entanglement of purification	30
5.1	Holographic entanglement of purification	30
5.2	Tensor networks from minimally entangled purifications	32
5.3	Multiple loops	36
6	Iteration and sub-AdS locality	38
6.1	The four-tensor network	38
6.2	Arbitrarily fine discretizations	41
7	Quantum geometry	43
7.1	Superpositions of tensor networks: sweeping away the “cobwebs”	43
7.2	Uncertainty relations for intersecting Ryu-Takayanagi surfaces	46
7.3	A no-go theorem for the four-tensor network	47
8	Discussion	53
8.1	Entanglement shadows	53
8.2	Bit threads	54
8.3	Fiber directions	55
8.4	Dynamics	56
8.5	Geometry from entanglement	57

1 Introduction

Among the most striking predictions of the AdS/CFT correspondence [1] is that the entanglement structure of a holographic CFT state is encoded in the geometry of its semiclassical gravitational dual. This conjectured correspondence is made precise for static spacetimes by the Ryu-Takayanagi (RT) formula [2, 3], and for dynamical spacetimes by the Hubeny-Rangamani-Takayanagi (HRT) formula [4]. These formulas relate the entanglement entropy of a CFT subregion to the area of an extremal surface in the bulk whose boundary coincides with that of the subregion, and were derived by path integral arguments in [5, 6].

This apparent holographic relationship between geometry and entanglement led to the proposal that tensor networks, which were originally developed as tools for the numerical analysis of condensed matter systems with restricted entanglement structure, might be a good toy model for the AdS/CFT correspondence [7, 8]. Tensor networks represent a quantum state on a D -dimensional lattice as a contraction of tensors lying on a $D + 1$ -dimensional graph with the lattice as its boundary, naturally disentangling the boundary state into a geometric “bulk” representation.

Since entanglement in the boundary state of a tensor network is related to the geometry of its bulk graph, tensor networks display holographic properties that are at least superficially similar to those of AdS/CFT. For example, all tensor networks follow a version of the Ryu-Takayanagi formula, in the sense that the entanglement entropy of a boundary subregion is bounded above by the “minimal area” bulk graph cut sharing a boundary with that subregion [7]. For a large class of tensor networks, this “Swingle bound” is exactly or approximately saturated [9, 10]. Many tensor networks also display features of quantum error correction [9, 11, 12], which are expected to appear in the AdS/CFT correspondence [13]. On the other hand, these models typically feature a flat (or almost flat) entanglement spectrum, which is at odds with known entanglement features of AdS/CFT (for more discussion on this point, see [10]).

In much of the existing literature, the idea that AdS/CFT can be explained in terms of tensor networks is taken “seriously, but not literally.” Holographic tensor networks are generally regarded as toy models for AdS/CFT that provide some intuition for how the geometric structure of a spacetime is encoded in the entanglement of its boundary dual. Even when a tensor network is interpreted literally as an approximate holographic description of a CFT state, it is often assumed that each tensor in the network must represent a volume of space that is at least of order ℓ_{AdS}^{d-1} (where ℓ_{AdS} is the characteristic scale of the semiclassical bulk spacetime, and d is its spacetime dimension).¹

This conclusion, however, is fundamentally at odds with the notion that tensor networks can be used to understand the Ryu-Takayanagi formula. Since the Ryu-Takayanagi formula is believed to hold exactly up to quantum and stringy corrections, it cannot be described adequately by a model that displays only AdS-scale locality. Any true tensor-network explanation of the Ryu-Takayanagi formula must therefore involve tensor network models that describe spacetime accurately at sub-AdS scales (so long as those scales remain above the string and Planck scales).

¹For a previous effort to construct tensor networks below the AdS scale, see [14].

The Ryu-Takayanagi formula is not the only geometric formula for an information-theoretic quantity that is claimed to hold below AdS scales. Progress in understanding the holographic properties of tensor networks has been accompanied in recent years by a series of increasingly bold conjectures regarding the formal relationship between entanglement and geometry in AdS/CFT (e.g., the holographic entanglement of purification conjecture [15, 16], which will play a central role in this paper). While many of these conjectures are partially inspired by tensor network models, they too are supposed to hold up to quantum (and stringy) corrections — well below the scale of existing tensor network models.

In this paper, we take literally the idea that there exists an approximate tensor network description for any holographic state, for essentially any discretization of the bulk, so long as the Planck and string scales are small compared to the discretization scale. As opposed to previous tensor network constructions, where a tensor network toy model is first defined and then shown to have properties similar to AdS/CFT, here we start with a holographic state in full AdS/CFT and construct a tensor network that describes it with high accuracy. Given a holographic CFT state with a static semiclassical dual, and assuming certain conjectures that naturally extend the holographic entanglement of purification conjecture, we provide an explicit procedure to construct a network that is “geometrically appropriate” for the holographic state in the following sense:

- (i) it approximately reproduces the original CFT state on its boundary, and
- (ii) it has the same geometric features as the bulk dual to leading order in N .

Our networks include subleading fluctuations around a flat entanglement spectrum, which we interpret as corresponding to fluctuations of the areas of extremal surfaces in full AdS/CFT. As a result, our constructions do not suffer from the usual issue of a flat, non-physical entanglement spectrum.²

If our assumptions hold up to quantum and stringy corrections (as is commonly assumed for the Ryu-Takayanagi formula and the holographic entanglement of purification conjecture), then essentially any discretization of a bulk geometry gives a corresponding tensor network description of the boundary state at sufficiently large N (and strong coupling). The tensor network description of the AdS/CFT correspondence would therefore be valid even at sub-AdS scales, and could be interpreted not as a “toy model” but as a genuine description of the quantum gravitational theory.

When the chosen discretization of the bulk geometry is constructed using only non-intersecting minimal surfaces, we are able to prove the existence of a corresponding tensor network, whose graph will always form a tree, without resorting to the holographic entanglement of purification. The construction relies only on the Ryu-Takayanagi formula for the von Neumann entropy and its extension to more general formulas for holographic Rényi entropies in [17]; it is only when we extend our construction to finer-grained discretizations (and hence tensor networks with loops) that we require the use of entropies of purification. Regardless of the particular discretization in question, all of our procedures involve

²Our tensor networks will not have the same Rényi entropies as the AdS/CFT states from which they are constructed, nor do they need to in order for condition (i) to be satisfied. As mentioned in section 3.1, the Rényi entropies are very sensitive to small changes in a state that do not alter its physical properties.

a process of disentangling and removing as many boundary degrees of freedom as possible to simplify the entanglement structure of a holographic state without altering its physical properties. In this sense, our constructions form a systematic approach for building tensor networks that describe their corresponding holographic states with maximal efficiency.

In section 2, we review the fundamental principles of tensor networks, introduce the abstract index notation that will be used in the remainder of the paper, and explain the basics of entanglement distillation from the tensor network perspective. We also sketch an important, but unpublished, result due to Hayden, Swingle, and Walter [18] on one-shot holographic entropies that allows us to apply the general entanglement distillation procedure in the holographic context. In section 3, we then identify a large class of tensor networks, which we call “tree tensor networks,” that can be constructed rigorously from holographic CFT states using one-shot entanglement distillation. In section 4, we show how these constructions can be adapted to produce holographic quantum error correcting codes. In section 5, we review the holographic entanglement of purification conjecture, and show that it can be used to improve the granularity of our networks by localizing the information contained within a single Ryu-Takayanagi surface. In section 6, we explain how natural generalizations of the holographic entanglement of purification conjecture can be used to extend this procedure to produce even finer-grained tensor networks, where the discretization scale of the network may lie well below the AdS scale so long as it exceeds the string and Planck lengths. In section 7, we discuss quantum fluctuations of the spacetime geometry and propose an interpretation in terms of quantum superpositions of tensor networks, which require significantly fewer degrees of freedom to describe than a full, “fluctuating” geometry. We identify an uncertainty relationship between the areas of intersecting Ryu-Takayanagi surfaces, and prove a related no-go theorem that limits the kinds of bulk-to-boundary isometries that can be obtained in a tensor network that accurately reproduces the geometry of AdS/CFT. Finally, in section 8, we summarize our essential results and present several potential avenues for future work.

2 Tensor networks, entanglement distillation, and one-shot quantum information

2.1 Tensor networks

A pure state $|\psi\rangle$ on a multipartite Hilbert space $\mathcal{H} = \mathcal{H}_{A_1} \otimes \cdots \otimes \mathcal{H}_{A_n}$ may be thought of as a tensor with n (abstract) up-indices, each one corresponding to a tensor factor of \mathcal{H} . In the tensor interpretation, we write such a state as

$$|\psi\rangle \leftrightarrow \psi^{A_1 \dots A_n}. \quad (2.1)$$

Such a tensor can generally be written as an outer product and contraction of other tensors, each of which acts only on some subset of the tensor factors A_1 through A_n . For $n = 2$, for example, a state might be written as

$$\psi^{A_1 A_2} = P^{A_1 B} Q^{A_2 C}{}_B, \quad (2.2)$$

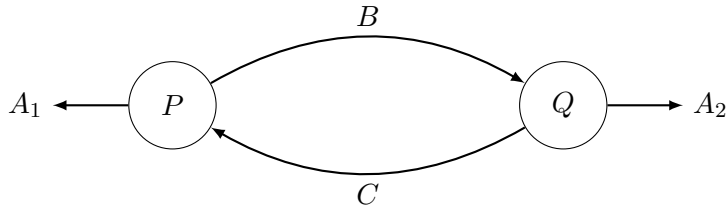


Figure 1. A simple network for the state given in equation (2.2). Outward-pointing arrows denote up-indices, inward-pointing arrows denote down-indices, and arrows connecting tensors denote contractions.

where the indices B and C correspond to some auxiliary Hilbert spaces \mathcal{H}_B and \mathcal{H}_C (and their dual spaces \mathcal{H}_B^* and \mathcal{H}_C^*), defined solely for the purpose of constructing tensors P and Q . Note that up-indices always refer to vector spaces, while down-indices refer to their corresponding dual spaces.

While it is always possible to find *some* outer product representation of any given tensor, there are significant computational advantages to finding one in which the contracted Hilbert spaces have small dimension compared to the physical Hilbert space factors. (These contracted Hilbert spaces are often referred to as “bonds,” with their dimensions referred to as the “bond dimensions” of the network.) Even in cases where the bond dimension is chosen to be of the same order as the physical Hilbert space dimension, many physically interesting states have outer product representations with some particular restricted structure that allows them to be simulated efficiently on a classical computer (see, e.g., [19] and [20]).

From the perspective of holography, one advantage of such an outer product representation of a quantum many-body state is that it has a natural geometric interpretation that shares features with holographic spacetimes. This geometric interpretation is called a *tensor network*. A tensor network is constructed from an outer product representation of a quantum state by drawing a vertex for each tensor, with one edge for each of its indices. In our convention, edges corresponding to up-indices will be labeled with arrows that point away from the vertex, while edges corresponding to down-indices will be labeled with arrows that point toward the vertex. Contractions are denoted by connecting the corresponding edges. As a simple example, the tensor network corresponding to the state in equation (2.2) is given in figure 1.

An important connection between tensor networks and quantum information theory arises from the fact that each bond in a tensor network can be thought of as a projection onto a maximally entangled state in the bond Hilbert space. For concreteness, consider the tensor network representation

$$\psi^{AB} = P^{A\gamma} Q^B_{\gamma} \tag{2.3}$$

of a state on the bipartite Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$. The tensor network is constructed by contraction over some bond Hilbert space \mathcal{H}_γ . The inner product on this bond Hilbert space selects a preferred maximally entangled state $|\phi\rangle$ on the product space $\mathcal{H}_\gamma \otimes \overline{\mathcal{H}}_\gamma$, where

$\overline{\mathcal{H}}_\gamma$ is the complex conjugate vector space to \mathcal{H}_γ .³ Since the state is maximally entangled, tracing out either of its tensor factors yields the identity operator on the remaining tensor factor (up to a normalization factor). In index notation, this statement is simply

$$\phi^{\gamma\bar{\gamma}}\phi_{\gamma'\bar{\gamma}}^* = \frac{1}{d}\delta^{\gamma}_{\gamma'}, \tag{2.4}$$

$$\phi^{\gamma\bar{\gamma}}\phi_{\gamma\bar{\gamma}'}^* = \frac{1}{d}\delta^{\bar{\gamma}}_{\bar{\gamma}'}, \tag{2.5}$$

where ϕ^* denotes the tensor corresponding to the dual state $\langle\phi|$. Equations (2.4) and (2.5) imply that ϕ and ϕ^* can be used to raise and lower indices between \mathcal{H}_γ and $\overline{\mathcal{H}}_\gamma$. Since ϕ^* can be obtained from ϕ by lowering its indices, we will generally drop the asterisk and refer to the tensors as $\phi^{\gamma\bar{\gamma}}$ and $\phi_{\gamma\bar{\gamma}}$, respectively.

Using ϕ to raise and lower indices, the state in equation (2.3) can be rewritten as

$$\psi^{AB} = P^{A\gamma}Q^{B\bar{\gamma}}\phi_{\gamma\bar{\gamma}}, \tag{2.6}$$

which is a slightly different tensor network representation of the same state. In the familiar bra-ket notation, this corresponds to projecting a state $|P\rangle \in \mathcal{H}_A \otimes \mathcal{H}_\gamma$ and a state $|Q\rangle \in \mathcal{H}_B \otimes \overline{\mathcal{H}}_\gamma$ onto a maximally entangled state $|\phi\rangle \in \mathcal{H}_\gamma \otimes \overline{\mathcal{H}}_\gamma$, i.e.,

$$|\psi\rangle = \langle\phi|(|P\rangle \otimes |Q\rangle). \tag{2.7}$$

A tensor network in which every bond takes this form is sometimes referred to as a *projected entangled-pair state* (PEPS) network (see, e.g., [10]), though the term PEPS is more commonly used to refer to a more highly-restricted class of tensor networks on a (usually square) lattice where each tensor has a single uncontracted physical index [21]. To avoid confusion, we refer to a tensor network with bonds of the form (2.7) as a *projection of entangled pairs* (PEP). Note that any tensor network can be rewritten in this form by raising and lowering indices with the appropriate maximally entangled state on each bond.

The original tensor network given by equation (2.3) is drawn in figure 2a, and its equivalent PEP network in figure 2b. Since PEP networks manifestly represent bonds as maximally entangled states, it is generally useful to consider them when drawing connections between tensor networks and AdS/CFT, where geometric features of a semiclassical spacetime correspond directly to entanglement features of its boundary dual.

Particular classes of tensor networks are known to reproduce various features of AdS/CFT [9, 10], such as a version of the Ryu-Takayanagi formula and holographic quantum error correction. In this paper, we work backwards, beginning with these known features of AdS/CFT and using them to *produce* tensor networks with corresponding properties. Much of our protocol is reliant on the procedure of *entanglement distillation*, which shows how the entanglement between subregions of some physical state can be distilled out of the state in the form of a large number of EPR pairs (which will, ultimately, become the maximally entangled bonds of a PEP-style network). This procedure is the subject of the following subsection.

³Formally, the inner product on \mathcal{H}_γ is a bilinear map $L : \mathcal{H}_\gamma \times \overline{\mathcal{H}}_\gamma \rightarrow \mathbb{C}$, or, equivalently, a tensor in the space $\mathcal{H}_\gamma^* \otimes \overline{\mathcal{H}}_\gamma^*$. Since the inner product is nondegenerate by assumption, it has an inverse tensor L^{-1} on $\mathcal{H}_\gamma \otimes \overline{\mathcal{H}}_\gamma$, which can be shown to be a maximally entangled state on the tensor product Hilbert space.

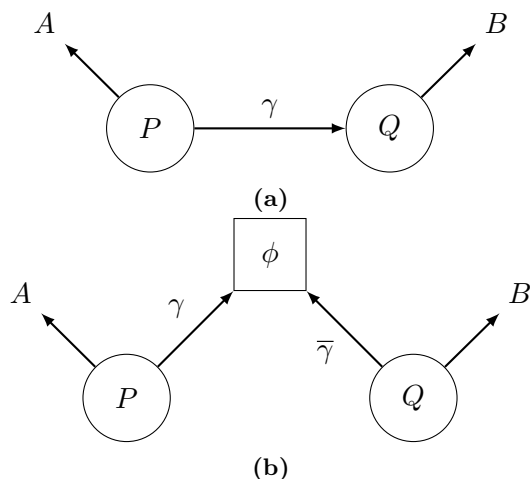


Figure 2. Tensor networks for equations (2.3) and (2.6). (a) A tensor network on a bipartite system with a single contraction. (b) A PEP network for the same state created by replacing the contraction with a maximally entangled state.

2.2 Entanglement distillation

Consider a state $|\psi\rangle$ in a CFT with large central charge that is known to have a static, semiclassical gravitational dual.⁴ If the domain of the CFT is partitioned into connected regions A and A^c , then the Ryu-Takayanagi formula states that the entanglement entropy of $|\psi\rangle$ between A and A^c is given to leading order in the gravitational constant G_N by the area of the minimal codimension-2 bulk surface anchored on ∂A and homologous to A , i.e.,

$$S(\psi_{(A)}) = S(\psi_{(A^c)}) = \min_{\gamma, \partial\gamma = \partial A} \frac{\text{area}(\gamma)}{4G_N} + O(1). \quad (2.8)$$

Implicit in this statement is a simultaneous regularization procedure where an ultraviolet cutoff is chosen in the CFT alongside a matching radial cutoff in the bulk spacetime [2].⁵

The Ryu-Takayanagi formula encourages us to think of the information encoded in the entanglement spectrum of A and A^c as lying physically on the extremal surface in the bulk that partitions A and A^c . For example, the Ryu-Takayanagi formula is sometimes interpreted as counting the number of “bit-threads” of entanglement, each of which occupy a Planckian area $1/4G_N$ of the extremal surface [24]. One way to more concretely justify this intuition comes from *entanglement distillation*, which makes precise the statement that entanglement entropy measures the number of qubits (or, more precisely, “ebits”) of entanglement shared between a region and its complement.

Entanglement distillation is the procedure by which a large number m of copies of some bipartite quantum state $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_{A^c}$ can be converted into a large number n of Bell

⁴In fact, it suffices to consider states whose semiclassical gravitational duals contain a moment of time reflection symmetry around the Cauchy slice being considered. This assumption is crucial to our construction, however, in section 8, we discuss the possibility of lifting this restriction in future work.

⁵That the subleading corrections to the Ryu-Takayanagi formula are $O(1)$ in G_N was established in [22] and [23].

pairs with some fixed asymptotic ratio $n/m \approx S(A)/\ln(2)$.⁶ For our purposes, the most useful formulation of this principle is that for large m , the state $|\psi\rangle^{\otimes m}$ can be expressed with high fidelity as

$$|\psi\rangle^{\otimes m} \approx (V \otimes W) \left[\frac{1}{\sqrt{D}} \sum_{i=0}^{e^{S(A)m-O(\sqrt{m})}} |i\bar{i}\rangle \otimes \sum_{j=0}^{e^{O(\sqrt{m})}} \sqrt{p_j} |j\bar{j}\rangle \right], \quad (2.9)$$

where V and W are isometries⁷ that embed Hilbert spaces of size $e^{S(A)m-O(\sqrt{m})}$ and $e^{O(\sqrt{m})}$, along with their complex conjugate Hilbert spaces, back into the physical space $\mathcal{H}_A \otimes \mathcal{H}_{A^c}$.⁸ (See, e.g., [25] for further details.)

The first factor of the tensor product in equation (2.9) is just a maximally entangled Bell state on $O(n)$ qubits, while the second term lives in a Hilbert space of subleading dimension in the asymptotic entanglement entropy $n \ln(2) = mS(A)$. The fact that asymptotically many copies of $|\psi\rangle$ can be approximately represented as a number of Bell pairs determined by the entanglement entropy gives partial justification for thinking of $S(A)$ as a measure of the number of degrees of freedom entangled between A and A^c , which in turn encourages us to think of all the information in the entanglement spectrum of a subregion of a holographic state as living physically on its Ryu-Takayanagi surface.

Using the notation of section 2.1, equation (2.9) is a tensor network of the form

$$\psi^{A(1)A^c(1)\dots A(m)A^c(m)} \approx V^{A(1)\dots A(m)} \gamma_f W^{A^c(1)\dots A^c(m)} \bar{\gamma}_{\bar{f}} \phi^{\gamma\bar{\gamma}} \sigma^{f\bar{f}}, \quad (2.10)$$

where $\phi^{\gamma\bar{\gamma}}$ is the maximally entangled state on $O(n)$ qubits and $\sigma^{f\bar{f}}$ is the leftover state on a Hilbert space of subleading dimension.

We see from equation (2.10) that entanglement distillation can be used to construct a simple tensor network that reproduces $|\psi\rangle^{\otimes m}$ with high fidelity for large m and any given quantum state $|\psi\rangle$. We could easily apply this procedure to a holographic CFT state (or, indeed, a non-holographic CFT state) and obtain a tensor network for the product state $|\psi^{\text{CFT}}\rangle^{\otimes m}$. In the holographic case, this tensor network will respect the geometry of AdS/CFT in the sense that its internal bond dimensions are inherited from the areas of the Ryu-Takayanagi surfaces.

Of course, in quantum gravity, people do not generally consider a large number of copies of a single holographic state. Any hope of understanding the entanglement structure of a *single* holographic state using the tools of quantum information theory is inhibited by the fact that almost all operational interpretations of a state's von Neumann entropy involve an asymptotic number of copies of the state in question. For our purposes, therefore, it

⁶In most of the original literature, “entanglement distillation” is used to refer to the general case of a mixed state $\rho \in S(\mathcal{H}_A \otimes \mathcal{H}_{A^c})$, while “entanglement concentration” is used in the special case where ρ is pure. Here, we use the terms interchangeably.

⁷In the literature, the word “isometry” is commonly used to refer to a Hilbert space map V satisfying $V^\dagger V = \mathbf{1}$. These maps are not generally isomorphisms in the mathematical sense, i.e., they are not invertible, unless the domain and target spaces have the same dimension.

⁸This expression follows from a procedure of smoothing and binning the entanglement spectrum of $|\psi\rangle^{\otimes m}$ in a way that is analogous to the procedure for a holographic state detailed in section 3.

is instead useful to consider “one-shot” or “smooth” entropies, which determine how well procedures like entanglement distillation can be carried out using only a single copy of a state. Luckily, the one-shot entropies of holographic states are highly constrained by the large central charge of the CFT, which is the focus of the following section.

2.3 Holographic one-shot entropies

If the von Neumann entropy only has a physically meaningful interpretation in the limit of asymptotically many copies of a state, then it might seem peculiar that the holographic dictionary relates the von Neumann entropy of a single reduced state to the (physically significant) area of a corresponding Ryu-Takayanagi surface. The resolution to this puzzle is quite simple: as we shall see, the semiclassical limit of large N (or, equivalently, $G_N \rightarrow 0$) in a holographic state fulfills the same information-theoretic purpose as the asymptotic limit of a large number of identical states in non-holographic quantum information theory.

Two important entropy measures for a quantum state ρ are the max-entropy

$$S_{\max} = \log(\text{rank}(\rho)) \tag{2.11}$$

and the min-entropy

$$S_{\min} = \log(\lambda_{\max}^{-1}(\rho)), \tag{2.12}$$

where $\lambda_{\max}(\rho)$ is the largest eigenvalue of ρ . These quantities agree with the von Neumann entropy for a state with a flat probability spectrum, but generically differ for arbitrary density matrices while satisfying $S_{\max} \geq S \geq S_{\min}$. The max- and min-entropies can be interpreted as Rényi entropies

$$S_{\alpha} = \frac{1}{1-\alpha} \log \text{Tr}(\rho^{\alpha}) \tag{2.13}$$

in the limits $\alpha \rightarrow 0$ and $\alpha \rightarrow \infty$ respectively. (Note that the von Neumann entropy S is given by the Rényi entropy in the limit $\alpha \rightarrow 1$.)

Since we are only interested in constructing tensor network states to within some small tolerance ε , it is more natural for our purposes (and in many similar situations) to consider the smooth max-entropy

$$S_{\max}^{\varepsilon} = \min_{\|\rho-\sigma\|_1 < \varepsilon} \log(\text{rank}(\sigma)) \tag{2.14}$$

and the smooth min-entropy

$$S_{\min}^{\varepsilon} = \max_{\|\rho-\sigma\|_1 < \varepsilon} \log(\lambda_{\max}^{-1}(\sigma)). \tag{2.15}$$

In other words, we consider the minimum max-entropy (respectively the maximum min-entropy) of any state σ lying within a ball of radius ε around the state ρ , where we have used the trace norm $\|\rho-\sigma\|_1 = \text{Tr} \sqrt{(\rho-\sigma)^{\dagger}(\rho-\sigma)}$ as a metric on the space of density matrices.

For holographic theories, the *smooth* min- and max-entropies are expected to agree with the von Neumann entropy to leading order in $1/G_N$. This was shown in [26] for single intervals in ground and thermal states of 1 + 1-dimensional holographic CFTs (where the density of states of the modular Hamiltonian may be computed explicitly), and in [18] for

arbitrary regions in holographic theories of arbitrary dimension.⁹ Somewhat counterintuitively, this generic behavior of the smooth min- and max-entropies follows from the fact that in holographic theories, the Rényi entropies are given to leading order by

$$S_\alpha = \frac{s_\alpha}{G_N}, \tag{2.16}$$

where s_α is independent of G_N but depends non-trivially on α and is related to the areas of surfaces in particular backreacted geometries [17, 28]. Since [18] remains unpublished, we include here a simplified version of the derivation of the smooth min- and max-entropies for holographic states, which follows on very general grounds from (2.16).¹⁰

Let $K = -\log(\rho)$ be the modular Hamiltonian corresponding to ρ . Let $\{E_i\}$ be the eigenvalues of K , with density of states defined by

$$D(E) \equiv \sum_i \delta(E - E_i). \tag{2.17}$$

Then the partition function

$$Z(\alpha) = \int_0^\infty dE D(E) e^{-\alpha E} \tag{2.18}$$

is related to the Rényi entropies by

$$e^{(1-\alpha)S_\alpha} = \text{Tr}(\rho^\alpha) = Z(\alpha). \tag{2.19}$$

Hence $e^{(1-\alpha)S_\alpha}$ is the Laplace transform of $D(E)$ and thus the density of states is given by

$$D(E) = IL(e^{(1-\alpha)S_\alpha})(E) = \int_C d\alpha e^{\alpha E} e^{(1-\alpha)S_\alpha}, \tag{2.20}$$

where $IL(e^{(1-\alpha)S_\alpha})$ is the inverse Laplace transform of $e^{(1-\alpha)S_\alpha}$ and C is a contour parallel to the imaginary axis with sufficiently large positive real part.¹¹ If S_α has the form given in (2.16), then $D(E)$ can be evaluated by a saddle point approximation for sufficiently small G_N . To leading order, we must therefore find

$$D(E) = e^{f(G_N E)/G_N + o(1/G_N)} \tag{2.21}$$

for some function $f(G_N E)$ that can be found by evaluating the exponent in (2.20) at the saddle point. If we substitute this expression for the density of states back into the expression given in equation (2.18), and substitute $E' = EG_N$, we find that the trace of ρ can be written as

$$Z(1) = \int dE D(E) e^{-E} = \int dE' e^{(f(E') - E')/G_N + o(1/G_N)}. \tag{2.22}$$

⁹For discussion of smooth max-entropies in general quantum field theories, see [27].

¹⁰While this article was in preparation, a pair of articles [29, 30] appeared with a similar analysis of the Rényi entropy spectrum; we will discuss their proposal that tensor networks correspond to area-eigenstates in more detail in section 7.1.

¹¹In this case, it suffices to take real part greater than or equal to one, and probably greater than or equal to zero. (Assuming the modular Hamiltonian has no maximum temperature state.)

Note that the coefficient of G_N that would appear from substituting E' for E in the measure of the integral has been absorbed into the subleading corrections of order $e^{o(1/G_N)}$.

By the same argument given above for the density of states, we find that the integral for $Z(1)$ will be dominated at small G_N by the leading saddle point E'_0 . We can therefore approximate the integral to within any arbitrarily small precision ε by integrating over a restricted range of eigenvalues of the modular Hamiltonian, constraining E to lie in the range

$$E'_0/G_N - O(\sqrt{\log(1/\varepsilon)/G_N}) < E < E'_0/G_N + O(\sqrt{\log(1/\varepsilon)/G_N}). \quad (2.23)$$

Since E'_0 is the solution to the saddle point equation $f'(E') = 1$, it is independent of G_N .

The error in approximating this integral controls the error induced by shaving off the largest and smallest eigenvalues of ρ . More precisely, if we define a “smoothed state” $\sigma = P\rho P/\text{Tr}(P\rho P)$, where P is the projector onto the eigenspaces of $K = -\log(\rho)$ with eigenvalues in the range given in (2.23), then it is clear that (i) σ lies within an $O(\varepsilon)$ -ball of ρ , and (ii) it has maximal and minimal eigenvalues given by

$$\lambda_{\max}(\sigma) = e^{-E'_0/G_N + O(1/\sqrt{G_N})}, \quad (2.24)$$

$$\lambda_{\min}(\sigma) = e^{-E'_0/G_N - O(1/\sqrt{G_N})}. \quad (2.25)$$

It follows immediately from the definitions given in equations (2.14) and (2.15) that the smooth min- and max-entropies of ρ agree with one another to leading order in G_N , since we have

$$\text{rank}(\sigma)\lambda_{\min}(\sigma) \leq \text{Tr}(\sigma) \leq \text{rank}(\sigma)\lambda_{\max}(\sigma). \quad (2.26)$$

More precisely, since σ is normalized with $\text{Tr}(\sigma) = 1$, equations (2.24) and (2.25) imply

$$S_{\max}(\sigma) = \log(\text{rank}(\sigma)) \leq \log\left(\frac{1}{\lambda_{\min}(\sigma)}\right) = \frac{E'_0}{G_N} + O\left(\frac{1}{\sqrt{G_N}}\right) \quad (2.27)$$

and

$$S_{\min}(\sigma) = \log(\lambda_{\max}^{-1}(\sigma)) = \frac{E'_0}{G_N} - O\left(\frac{1}{\sqrt{G_N}}\right). \quad (2.28)$$

Since $S_{\max}(\sigma)$ must be greater than $S_{\min}(\sigma)$, equations (2.27) and (2.28) together imply that the min and max entropies of σ agree with one another to leading order in G_N . Since σ is within an $O(\varepsilon)$ -ball of ρ , the same statement holds true for the *smooth* min and max entropies of ρ , as we previously claimed.

The only remaining question is to find the saddle point value E'_0 . The von Neumann entropy of ρ is given by

$$S(\rho) = \int dE D(E) E e^{-E} = \frac{E'_0}{G_N} + O(1), \quad (2.29)$$

where we have used the form of $D(E)$ given in (2.17). It follows from the Ryu-Takayanagi formula that the saddle point value E'_0 is given by $A/4$, where A is the area of the corresponding RT surface. The smooth min- and max-entropies are therefore equal to the von Neumann entropy up to $O(1/\sqrt{G_N})$ corrections for any fixed nonzero ε . Since the von Neumann entropy S is of order $O(1/G_N)$, the corrections grow as the square root of

the entropy (for fixed values of the UV cutoff). It follows that for holographic states, the smooth min and max entropies satisfy

$$S_{\min} = S - O(\sqrt{S}), \tag{2.30}$$

$$S_{\max} = S + O(\sqrt{S}). \tag{2.31}$$

This is exactly the same scaling that is seen when we take the asymptotic limit of a large number of copies of a state.¹² The semiclassical holographic limit of large central charge is therefore replicating, at least partially, the effects of the asymptotic i.i.d. limit of a large number of identical copies of a single, non-holographic state. When we construct holographic tensor networks in the following section, we will see that these subleading corrections to the smooth min- and max-entropies can be related to known subleading contributions to holographic entanglement in AdS/CFT.

3 Tree tensor networks from holographic entanglement distillation

In order to construct a meaningful tensor network for a state in the AdS/CFT correspondence, it is necessary to produce a network that (i) reproduces the correct boundary state with high fidelity, and (ii) has a bulk geometry that matches the bulk spacetime. Since tensor networks have discrete geometries, property (ii) must be interpreted in terms of some discretization of the bulk spacetime. In this section, we consider *tree networks* — those constructed by discretizing the bulk with non-intersecting Ryu-Takayanagi surfaces. Given such a discretization, the underlying graph of the corresponding tensor network is taken to be the dual graph of the set of Ryu-Takayanagi surfaces and their corresponding boundary regions. A sample discretization of vacuum AdS_3 , along with the corresponding dual graph, is shown in figure 3.

Once a network is constructed on this graph, it is straightforward to quantify how well the resulting state satisfies property (i) by looking at the inner product between the state constructed by the tensor network and the target state. We will say that the network satisfies property (ii) if the bond dimension of each edge in the network matches the area of the Ryu-Takayanagi surface through which it passes. More precisely, we will require that each bond γ in the network satisfies $\dim(\mathcal{H}_\gamma) = e^{A_\gamma/4G_N + o(1/G_N)}$, where A_γ is the area of the corresponding RT surface.¹³ Since the dimension of any bond Hilbert space could be made arbitrarily large without altering the state by adding zero probability states, we also require that each contraction is *full rank in the bond space* in the sense that it contains no trivial

¹²Until now, we have assumed that ε is some fixed small number that is independent of G_N . However, the range of integration required to approximate (2.22) depends only weakly on the allowed error ε as $\sqrt{\log(1/\varepsilon)}$. Hence we can make the error ε non-perturbatively small with respect to G_N , at the small cost of allowing the smooth min- and max-entropies be separated by $O(\sqrt{Sf(S)})$ for some super-logarithmic function $f(S)$. (We can then obtain $\varepsilon = e^{-f(S)}$, which is non-perturbatively small.) A natural choice might be $f(S) = S^\delta$ for some small δ , or maybe $f(S) = (\log(S))^2$. Regardless of the specific function chosen, it is easy to ensure that $O(\sqrt{Sf(S)})$ is subleading compared to the entropy S .

¹³Both sides of this equality are infinite. As usual, equations involving the entanglement entropy or the area of extremal surfaces should be interpreted in the context of a regularization scheme in which the CFT state is regulated on a lattice with finite spacing and the bulk spacetime is regulated with a radial cutoff.

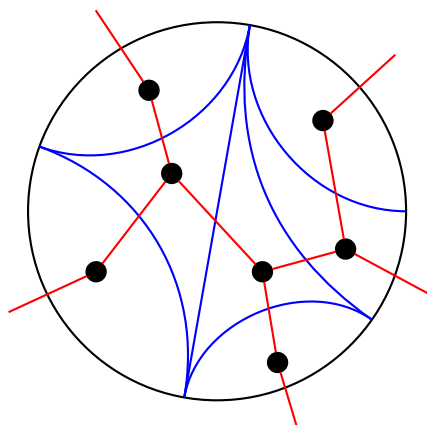


Figure 3. A bulk discretization of vacuum AdS_3 by non-intersecting RT surfaces. Blue curves represent extremal surfaces in the bulk, red lines are edges in the dual graph, and black dots are vertices in the dual graph. In the resulting tensor network, the dangling edges passing through the boundary of the spacetime will correspond to uncontracted (physical) Hilbert space indices.

contractions.¹⁴ In fact, this condition is not quite strong enough, as the dimension of \mathcal{H}_γ could still be made arbitrarily large by the addition of states with arbitrarily small probability — in our construction, this possibility is avoided by ensuring that the eigenvalues of each bond are bounded below by a nontrivial function of the bond entropy (cf. the smoothing process of section 2.3, where each “smoothed state” has bounded minimal eigenvalue).

3.1 Bipartite tensor networks

To explain the procedure of constructing tree networks via one-shot entanglement distillation, we restrict temporarily to the case where the bulk is discretized by a single Ryu-Takayanagi surface. For such a discretization, the boundary is partitioned into two connected regions A and A^c . The generic procedure for arbitrary non-intersecting bulk partitions is detailed in section 3.2.

As discussed in section 2.3, the smooth min- and max-entropies of a holographic state agree with the von Neumann entropy to leading order in G_N . It follows that for the reduced CFT density matrix $\psi_{(A)}$ of a holographic CFT state $|\psi\rangle$, there exists a normalized state $\psi_{(A)}^\varepsilon$ within ε trace distance of $\psi_{(A)}$ satisfying

$$\text{rank} \left(\psi_{(A)}^\varepsilon \right) = e^{S(A)+O(\sqrt{S})}, \tag{3.1}$$

$$\lambda_{\max} \left(\psi_{(A)}^\varepsilon \right) = e^{-S(A)+O(\sqrt{S})}, \tag{3.2}$$

where $\lambda_{\max} \left(\psi_{(A)}^\varepsilon \right)$ is the largest eigenvalue of $\psi_{(A)}^\varepsilon$.

¹⁴Formally, for a bond of the form $P^{A\gamma}Q^B_\gamma$, we require that there is no nonzero vector v^γ in \mathcal{H}_γ or dual vector ω_γ in \mathcal{H}_γ^* such that $P^{A\gamma}\omega_\gamma$ or $Q^B_\gamma v^\gamma$ identically vanishes.

Given the full, pure CFT state $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_{A^c}$,¹⁵ one can write the Schmidt decomposition

$$|\psi\rangle = \sum_n \sqrt{\lambda_n} |n\rangle_A |n\rangle_{A^c}, \quad (3.3)$$

where $\{\lambda_n\}$ are the eigenvalues of the reduced states $\psi_{(A)}$ and $\psi_{(A^c)}$. If $\{\tilde{\lambda}_n\}$ are the eigenvalues of the smoothed state $\psi_{(A)}^\varepsilon$, then it is easy to verify that the state

$$|\psi^\varepsilon\rangle = \sum_n \sqrt{\tilde{\lambda}_n} |n\rangle_A |n\rangle_{A^c} \quad (3.4)$$

approximates the original state $|\psi\rangle$ with very high fidelity. In particular, we have

$$|\langle \psi^\varepsilon | \psi \rangle|^2 = F(\psi_{(A)}, \psi_{(A)}^\varepsilon) \geq \left(1 - \frac{1}{2} \|\psi_{(A)} - \psi_{(A)}^\varepsilon\|_1\right)^2 \geq 1 - \varepsilon, \quad (3.5)$$

where $F(\rho, \sigma) = [\text{Tr} \sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}]^2$ is the *fidelity* of two quantum states. The first equality in (3.5) follows from the definition of fidelity and the form of the states (3.3) and (3.4), while the subsequent inequality is one of the Fuchs-van de Graaf inequalities [31].

If we re-order the probability spectrum $\tilde{\lambda}_n$ such that it is monotonically decreasing, i.e. $\tilde{\lambda}_{n+1} \leq \tilde{\lambda}_n$, and break the resulting sum into blocks of size Δ , then we may rewrite (3.4) as

$$|\psi^\varepsilon\rangle = \sum_{n=0}^{\text{rank}[\psi_{(A)}^\varepsilon]/\Delta-1} \sum_{m=0}^{\Delta-1} \sqrt{\tilde{\lambda}_{n\Delta+m}} |n\Delta+m\rangle_A |n\Delta+m\rangle_{A^c} \quad (3.6)$$

Now, suppose we discard the m -dependence of the eigenvalues $\tilde{\lambda}_{n\Delta+m}$ and replace all of the eigenvalues in each block with the average value of that block, $\tilde{\lambda}_{n\Delta}^{\text{avg}}$. The resulting state (which is still correctly normalized) is

$$|\Psi^\varepsilon\rangle = \sum_{n=0}^{\text{rank}[\psi_{(A)}^\varepsilon]/\Delta-1} \sum_{m=0}^{\Delta-1} \sqrt{\tilde{\lambda}_{n\Delta}^{\text{avg}}} |n\Delta+m\rangle_A |n\Delta+m\rangle_{A^c}, \quad (3.7)$$

and satisfies

$$\|\Psi^\varepsilon - \psi_{(A)}^\varepsilon\|_1 \leq \lambda_{\max}[\psi_{(A)}^\varepsilon] \cdot \Delta \equiv \delta. \quad (3.8)$$

By the same arguments as in (3.5), the overlap between $|\Psi^\varepsilon\rangle$ and the original CFT state $|\psi\rangle$ is bounded below by

$$|\langle \psi | \Psi^\varepsilon \rangle|^2 \geq |\langle \psi | \psi^\varepsilon \rangle|^2 |\langle \psi^\varepsilon | \Psi^\varepsilon \rangle|^2 \geq 1 - \varepsilon - \delta. \quad (3.9)$$

If we choose $\Delta = e^{S-O(\sqrt{S})}$, with the $O(\sqrt{S})$ dependence chosen to approximately cancel the $O(\sqrt{S})$ dependence of $\lambda_{\max}[\psi_{(A)}^\varepsilon]$ given in equation (3.2), then we can ensure that δ remains small, while the state becomes

$$|\Psi^\varepsilon\rangle = \sum_{n=0}^{e^{O(\sqrt{S})}} \sum_{m=0}^{e^{S-O(\sqrt{S})}} \sqrt{\tilde{\lambda}_{n\Delta}^{\text{avg}}} |n\Delta+m\rangle_A |n\Delta+m\rangle_{A^c}, \quad (3.10)$$

¹⁵In reality, the Hilbert space of the actual CFT will not factorize in this way, due to ultraviolet issues. However, since we have already regularized the theory, there is no problem splitting the Hilbert space into tensor factors.

To properly distill the EPR pairs out of this state, we define auxiliary Hilbert spaces \mathcal{H}_γ and \mathcal{H}_f with dimensions given by

$$\dim \mathcal{H}_f = e^{O(\sqrt{S})}, \quad (3.11)$$

$$\dim \mathcal{H}_\gamma = e^{S-O(\sqrt{S})}, \quad (3.12)$$

where the precise values are chosen to match the range of the sums in equation (3.10). We define the isometries $\mathcal{H}_f \otimes \mathcal{H}_\gamma \hookrightarrow \mathcal{H}_A$ and $\overline{\mathcal{H}}_f \otimes \overline{\mathcal{H}}_\gamma \hookrightarrow \mathcal{H}_{A^c}$ by

$$V|n\rangle_f|m\rangle_\gamma = |n\Delta + m\rangle_A, \quad (3.13)$$

$$W|\overline{n}\rangle_{\overline{f}}|\overline{m}\rangle_{\overline{\gamma}} = |n\Delta + m\rangle_{A^c} \quad (3.14)$$

for some arbitrarily chosen bases of the auxiliary Hilbert spaces and the corresponding bases in their complex conjugate Hilbert spaces. We may then rewrite the state $|\Psi^\varepsilon\rangle$ as

$$|\Psi^\varepsilon\rangle = (V \otimes W) \left[\sum_{n=0}^{e^{O(\sqrt{S})}} \sqrt{\tilde{\lambda}_{n\Delta}^{\text{avg}}} |n\overline{n}\rangle_{f\overline{f}} \right] \otimes \left[\sum_{m=0}^{e^{S-O(\sqrt{S})}} |m\overline{m}\rangle_{\gamma\overline{\gamma}} \right]. \quad (3.15)$$

This expression is essentially identical to (2.9), except that it approximates a *single* copy of the original CFT state $|\psi\rangle$.

Approximate states such as those given in equation (3.15) are considerably more common in the quantum information literature than in discussions of quantum gravity, so we pause briefly to discuss their physical relevance. The first and most important point to note is that expectation value of any bounded operator \hat{O} on the Hilbert space of $|\psi\rangle$ is well approximated by its expectation value in the approximate state. In particular, since the overlap between the original CFT state $|\psi\rangle$ and the new state $|\Psi^\varepsilon\rangle$ takes the form given in (3.9), we can guarantee that the expectation values of bounded operators between the two states differ at most by

$$\left| \langle \psi | \hat{O} | \psi \rangle - \langle \Psi^\varepsilon | \hat{O} | \Psi^\varepsilon \rangle \right| \leq 2\sqrt{\varepsilon + \delta} \|\hat{O}\|. \quad (3.16)$$

If both ε and δ can be made sufficiently small, then this bound is quite narrow. Since the time evolution operator e^{iHt} is bounded above by $\|e^{iHt}\| \leq 1$ for any t that is either real or in the upper half-plane, a similar bound holds for correlation functions at arbitrary times. In particular, correlation functions of bounded operators satisfy

$$\left| \langle \psi | \hat{O}_1(t_1) \dots \hat{O}_n(t_n) | \psi \rangle - \langle \Psi^\varepsilon | \hat{O}_1(t_1) \dots \hat{O}_n(t_n) | \Psi^\varepsilon \rangle \right| \leq 2\sqrt{\varepsilon + \delta} \|\hat{O}_1\| \dots \|\hat{O}_n\|. \quad (3.17)$$

The states $|\psi\rangle$ and $|\Psi^\varepsilon\rangle$ therefore generally produce approximately the same values for any Euclidean or Lorentzian correlation function with arbitrarily large time gaps, including out-of-time order Lorentzian correlation functions.

There are a few scenarios in which $|\psi\rangle$ and $|\Psi^\varepsilon\rangle$ can display qualitatively different behavior; we argue that none of these scenarios are actually physically important. If $\langle \psi | \hat{O} | \psi \rangle$ is itself very small compared to $\|\hat{O}\|$, then there may $O(1)$ differences in the

relative size of $\langle \psi | \hat{O} | \psi \rangle$ and $\langle \Psi^\varepsilon | \hat{O} | \Psi^\varepsilon \rangle$. However, in this case, since ε can be arbitrarily small, both the expectation values would have to be zero at leading order; they would still agree up to $O(\varepsilon)$ corrections.

Secondly, non-observable quantities such as Rényi entropies for $\alpha \neq 1$ may (and will) look very different for $|\psi\rangle$ and $|\Psi^\varepsilon\rangle$.¹⁶ This is unsurprising, as the Rényi entropies are very sensitive to small and (physically) insignificant perturbations, and can vary drastically without significantly altering the expectation values of bounded operators.

Finally, we are often interested in unbounded operators. In this case we do not have any bound on the error in expectation values for the approximate state. However, in such circumstances we can generally replace an unbounded self-adjoint operator \hat{O} by some bounded function $f(\hat{O})$ of the operator without affecting the important physics. We will then obtain a tight bound on the error of the expectation value $\langle f(\hat{O}) \rangle$ when approximating it in the distilled state $|\Psi^\varepsilon\rangle$.

So long as we originally chose ε small and chose Δ correctly to ensure that δ is small, then (3.9) ensures that the distilled state $|\Psi^\varepsilon\rangle$ is a good approximation of the original CFT state $|\psi\rangle$. As discussed above, this implies that the physics of the two states should be the same up to non-perturbatively small corrections. Moreover, the expression given in (3.15) is a tensor network with a geometry matching the semiclassical dual of $|\psi\rangle$. In the abstract index notation of section 2.1, the state in (3.15) is written as

$$\Psi^{AA^c} = V^A_{f\gamma} W^{A^c}_{\bar{f}\bar{\gamma}} \phi^{\gamma\bar{\gamma}} \sigma^{f\bar{f}}, \tag{3.18}$$

where

$$|\phi\rangle = \sum_{m=0}^{e^{S-O(\sqrt{S})}} |m\bar{m}\rangle_{\gamma\bar{\gamma}}, \tag{3.19}$$

$$|\sigma\rangle = \sum_{n=0}^{e^{O(\sqrt{S})}} \sqrt{\tilde{\lambda}_{n\Delta}^{\text{avg}}} |n\bar{n}\rangle_{f\bar{f}}. \tag{3.20}$$

The network corresponding to (3.18) is sketched in figure 4, superposed over the corresponding discretization of vacuum AdS_3 . The tensors V and W correspond to the entanglement wedges of regions A and A^c , respectively, as isometries that embed the states ϕ and σ into the boundary. Since ϕ is a maximally entangled state on a Hilbert space of dimension $e^{O(S(A))}$, it has the right entanglement to reproduce the Ryu-Takayanagi surface that separates A from A^c . Because the bond dimension of the legs of the state $|\sigma\rangle$ is very small compared to the large bond dimensions of the ϕ -legs, we can think of the σ -legs as thin “cobwebs”, attached to the thick “girders” of the ϕ -legs.

It is important to be clear about our motivation in including a second approximation step, where we average $\tilde{\lambda}_i$ within each block and hence extract the dependence of $\tilde{\lambda}_i$ on i into the state $|\sigma\rangle$ in the (relatively) small auxiliary Hilbert space $\mathcal{H}_f \otimes \bar{\mathcal{H}}_f$. After all, we

¹⁶Here, we mean that the Rényi entropies are non-observable when provided with only a single copy of the state — they can be computed from matrix moments of a state, which are observable given multiple replicas of the state in question.

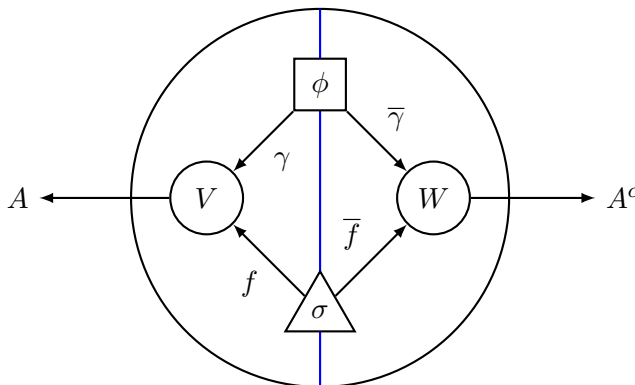


Figure 4. A tensor network for a bipartite discretization of AdS_3 by a single Ryu-Takayanagi surface. $\phi^{\gamma\bar{\gamma}}$ is a maximally entangled state on a Hilbert space of dimension $e^{O(S(A))}$, while $\sigma^{f\bar{f}}$ is a (generally not maximally entangled) state on a Hilbert space of dimension $e^{O(\sqrt{S(A)})}$. The tensors V and W embed these states isometrically into the boundary.

could already have constructed a tensor network with the correct discretized geometry by using the smoothed, but unflattened state

$$\sum_i \sqrt{\tilde{\lambda}_i} |i\rangle_\gamma |i\rangle_{\bar{\gamma}}. \quad (3.21)$$

Our purpose in flattening the $\tilde{\lambda}_i$ spectrum is not to claim that the state produced is in any sense “better” or “more holographic” than the smooth state that was constructed prior to this flattening procedure. In fact, the degrees of freedom extracted into $|\sigma\rangle$ are highly non-unique, as they depend a great deal on the exact block size Δ chosen while flattening the spectrum. However, by showing that the original state $|\psi\rangle$ can be approximated in this way, we are showing explicitly that the entanglement spectrum is so flat (up to smoothing) that the *effective* number of degrees of freedom in $|\psi\rangle$ that describe the gradient of the entanglement spectrum is subleading compared to the effective number of degrees of freedom that are simply maximally entangled between the two sides.

Note that to obtain a tensor network description with the correct leading order bond dimensions, we only needed the smooth max-entropy to be sufficiently small. The additional requirement that the smooth min-entropy agree with smooth max-entropy to leading order was imposed to ensure that the auxiliary Hilbert space \mathcal{H}_f on which the state is not maximally entangled has subleading dimension $e^{O(\sqrt{S})}$.

Before moving on to more general tree tensor networks, we pause to consider the role of the “cobweb” state $|\sigma\rangle$. This state does not have an immediate interpretation in the usual Ryu-Takayanagi picture of bulk entanglement, at least at leading order in G_N . Since the state $|\sigma\rangle$ is of subleading size in the entanglement entropy and hence in the central charge, the most natural interpretation is that it arises from quantum fluctuations in the bulk geometry (e.g., graviton fluctuations) that alter the areas of Ryu-Takayanagi surfaces

at subleading order in G_N .¹⁷ Generally, such fluctuations are expected to be suppressed by a factor of $O(\sqrt{G_N})$, meaning that the resulting fluctuations in the entropy are also of order $O(1/\sqrt{G_N})$.¹⁸ Since this matches the rank of $|\sigma\rangle$, it seems natural to associate the subleading state with these geometric fluctuations. We discuss this proposal in more detail in section 7.1.

It is worth commenting that we are also at liberty to absorb the state $|\sigma\rangle$ into one of (or a combination of) the isometries V and W . This simplifies the picture of the tensor network, but comes at the cost of at least one of the operators V and W no longer being an isometry. In fact, they will not even be approximate isometries, although they will remain isometries if interpreted as operators $V : \mathcal{H}_\gamma \hookrightarrow \overline{\mathcal{H}}_f \otimes \mathcal{H}_A$ and $W : \overline{\mathcal{H}}_\gamma \hookrightarrow \mathcal{H}_f \otimes \mathcal{H}_{Ac}$. The question of whether tensors in the network are (at least approximate) isometries is important for various reasons, both in ensuring that the boundary state of the network correctly approximates the original CFT state and in understanding the error correcting properties of the network. As such we shall always keep the state $|\sigma\rangle$, and its generalizations in more complicated networks, explicit.

3.2 General tree networks

The argument given above can be extended to construct a tree tensor network for an arbitrary discretization of the bulk by (non-intersecting) Ryu-Takayanagi surfaces. This generalization works roughly as one would expect: one simply localizes degrees of freedom to each RT surface in turn, each time creating an additional link and tensor in the network. However, some important subtleties arise during this process. It is easy to construct a superficially-reasonable procedure that will not actually approximate the original state with high fidelity. We will therefore describe the procedure for constructing generic tree networks in some detail.¹⁹

Our argument is inductive: we assume that we have successfully constructed a tree tensor network for a simpler discretization with one fewer RT surface, and then show that we can always add an additional RT surface while (approximately) preserving the bulk bond dimensions and the boundary CFT state. After arbitrarily many inductive steps, the final network will still approximate the original “target” CFT state on its boundary. However, to obtain rigorous bounds on our final error in approximating the original CFT state, we must be somewhat careful in the order in which we choose to add RT surfaces.

¹⁷An alternative approach [32, 33] describes the fluctuations of the areas of extremal surfaces as “edge modes,” i.e. superselection sectors that commute with the algebra of observables on both sides of the surface. In our approach, however, these fluctuations are described explicitly by the states of the $|\sigma\rangle$ tensor that lies on the Ryu-Takayanagi surface.

¹⁸To see that this is the correct scaling of metric fluctuations, note that a one graviton state with order-unity frequency has an $O(1/G_N)$ energy, but the energy is the square of the amplitude of the metric strain. Note however that in contexts where we are only interested in the *average* area, the $O(1/\sqrt{G_N})$ term does not appear because, for linearized gravitons, positive fluctuations are just as likely as negative fluctuations. That is why the quantum corrections in the holographic von Neumann entropy (2.8), which traces over the whole probability distribution, are merely $O(1)$. In our tensor network contexts, however, we need to keep track of the Hilbert space dimensions, which do not average out.

¹⁹For similar work in other contexts, see [34] and [35].

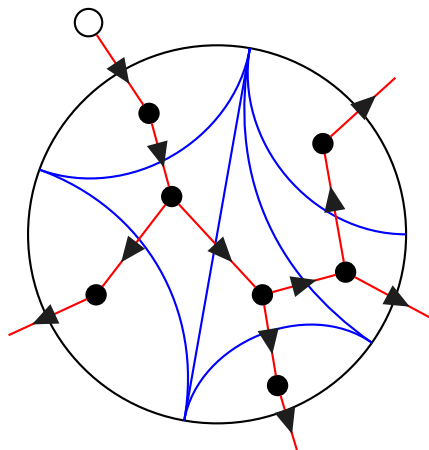


Figure 5. The bulk discretization of vacuum AdS_3 that was originally sketched in figure 3 has here been given a root-leaf orientation on its dual graph by choosing an arbitrary boundary edge as the “root” (represented here by a white circle).

Specifically, there must exist some choice of boundary node, which we shall label the “root” node, such that each “parent” tensor was added after all of its “children.” (In practice, it seems likely that one will obtain a correct approximation of the original state even when the RT surfaces are added in an arbitrary order. Without adding them according to a particularly nice orientation, however, it is hard to guarantee that one couldn’t obtain a large boundary error by sheer bad luck.)

We begin by choosing a “target” discretization of the bulk by non-intersecting RT surfaces, such as the one sketched above in figure 3. This will be the graph of our final tree tensor network. Designating one of its nodes as a root picks out a preferred orientation for the tree by flowing away from the root, as sketched in figure 5. In general, we will choose the root node to lie on the boundary, although our construction works even if the root is chosen to lie in the bulk. All boundary nodes that are not the root are now leaves of this oriented graph. Note that this orientation is defined only for the purpose of ordering the RT surfaces that make up the discretization, and is independent of the orientation imposed on the tensor network to denote up- and down-indices (cf. section 2.1).

To construct a tree tensor network for this graph, edges will be added to the network inductively from leaves up to the root. To preserve the isometry properties of the tensor network, no RT surface can be added to the network before all of its children have been added (according to the orientation induced by choosing a boundary root). Different choices of root node on the boundary, and even different orderings of RT surfaces that are consistent with a single root-leaf orientation, will in general produce different tree tensor networks. However, all such networks are geometrically appropriate for the AdS/CFT correspondence in the sense that they have bond dimensions that match the holographic geometry and boundary states that approximately reproduce the original CFT state.

To define the isometry properties of a tree tensor network precisely, we first define the state associated to any bulk region bounded by a mixture of Ryu-Takayanagi surfaces and

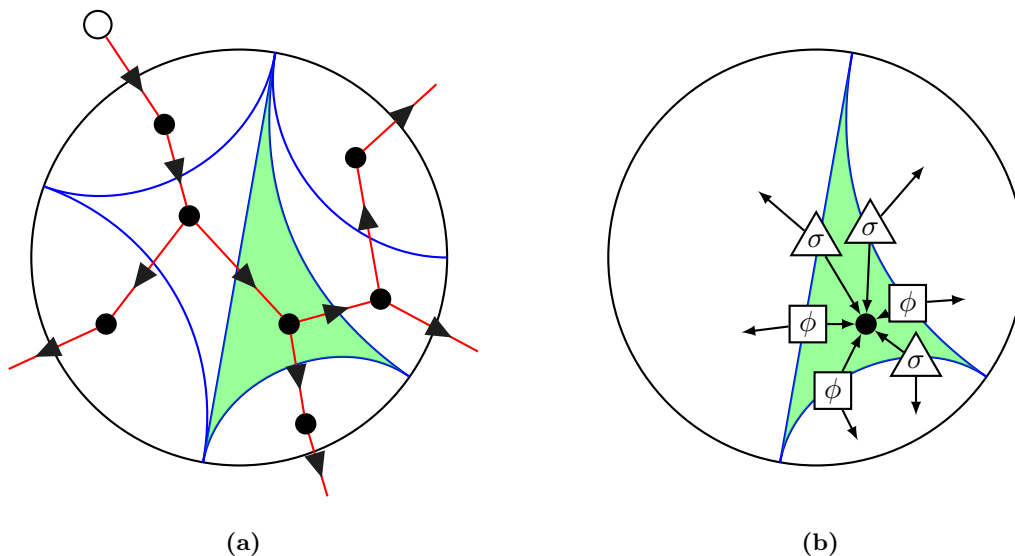


Figure 6. (a) The bulk discretization of vacuum AdS_3 shown in figure 5 can be divided into bulk states by selecting regions of the bulk that are bounded by Ryu-Takayanagi surfaces and boundary subregions. Such a region is shaded here. (b) The bulk state obtained from truncating a tree tensor network to the shaded region. Each edge of the tree tensor network is composed of a maximally entangled state $|\phi\rangle$ and a subleading state $|\sigma\rangle$, just as in the bipartite construction of section 3.1. The bulk state on the shaded region is defined by removing all tensors outside of the shaded region while keeping the edge states $|\phi\rangle$ and $|\sigma\rangle$ that define the edges at the boundary of the shaded region. The result is a truncated state on the tensor product Hilbert space of the edges that cross the boundary of the shaded region.

subregions of the boundary to be the state produced by a truncation of the tree tensor network to that region. This is sketched in figure 6 for a subregion of the tree tensor network that was introduced in figures 3 and 5. Importantly, the edge states $|\phi\rangle$ and $|\sigma\rangle$ associated to each RT surface are included in the state assigned to the region. There is a natural map from the state associated to a bulk region to the state associated to any larger region which contains the smaller region. This is essentially the “inclusion” map of the tensor network, which consists of all tensors that are included in the state of the larger region but not included in the smaller one; we call this the *extension map*. The extension map does not include the edge states associated to its “input” RT surfaces, as those are already included in the state of the smaller subregion; however, it will include the edge states associated to any “output” RT surfaces that bound the larger bulk subregion. This convention is chosen so that an extension map “beginning” on a given RT surface can always be composed with an extension map that “ends” on the same surface.

In the bipartite construction of section 3.1, the bulk tensors V and W each served as the extension map from the state in the complementary bulk region out to the global boundary. In this construction, the maps V and W were exact isometries. For a tree network constructed from a boundary root orientation (as sketched in figure 5), we will show inductively that extension maps flowing entirely along the direction of the orientation are always exact isometries. An extension map which flows partially against the orientation

of the graph will not in general be an exact isometry; however, it will be an approximate isometry with respect to a particular state-dependent metric. (For example, in figure 6a, the extension map from the shaded region through the network out to the “right half” of the boundary is an exact isometry, as it flows along the orientation of the graph. The extension map from the shaded region out to the global boundary, however, is only an approximate isometry, as it must flow against the orientation of the graph to reach the root node at the boundary.)

We now present the inductive argument for constructing a tree tensor network for an arbitrary (non-intersecting) bulk discretization. For clarity, this whole procedure is sketched in figure 7 for the final step of the oriented discretization sketched in figure 5. After a boundary edge has been designated as the root and a root-leaf orientation has been imposed on the dual graph, we pick one of the “uppermost” Ryu-Takayanagi surfaces (i.e., one of the surfaces that neighbors the root node) to be the last surface added to the network, and assume that we have already constructed a tree tensor network for the discretization that includes all but this final surface.²⁰ To make an inductive argument, we assume that the tensor network for the “all-but-one” discretization has been constructed so that it:

- (a) approximately reproduces the original “target” CFT state on the boundary,
- (b) has internal bond dimensions that match the areas of the discretization surfaces, and
- (c) has the isometry properties detailed above (i.e., extension maps that follow the flow of the root-leaf orientation are exact isometries).

We now consider the smallest bulk subregion a containing the new RT surface in its interior (cf. figure 7a, where the bulk subregion a is shaded). In the tensor network for the reduced discretization, the state associated to this region is formed by a single tensor together with $|\phi\rangle$ and $|\sigma\rangle$ states on each already-constructed RT surface on the boundary of a . Because we chose the RT surface to neighbor the root node, the extension map from this state to entire boundary state flows entirely along the orientation of the tree and so is an exact isometry by assumption. Furthermore, we have assumed that the entire boundary state is approximately equal to the target CFT state.

Since the smooth min- and max-entropies of the target state depend only weakly on the error ε , and since the extension map from a to the global boundary is an isometry, the smooth min- and max-entropies of the subregion state on a will agree with those of the target state to leading order. We can therefore apply the exact same bipartite distillation procedure to the subregion state on a that we used to construct the global bipartite tensor network in section 3.1. This will produce a bipartite tensor network $(V \otimes W)|\phi\rangle|\sigma\rangle$ of the form given in (3.15) that approximates the state of the subregion a . We label our isometries so that W is the “upwards” isometry that includes the root node in its image, while V is

²⁰The choice of “uppermost” RT surface is in general non-unique, and adding surfaces to the network in different orders will produce different tensor networks that all approximate the original CFT state with high fidelity.

the “downwards” isometry that maps away from the root node of the discretization (cf. figure 7c).

This newly-distilled state approximately reproduces the subregion state on a , and could be substituted directly into the network as in figure 7c. Doing so, however, would require erasing all of the $|\phi\rangle$ and $|\sigma\rangle$ states associated to each of the neighboring, previously-added RT surfaces, replacing them with the outward-pointing legs of V . Instead, we wish only to replace the central tensor associated to a . Fortunately, as discussed in section 2.1, there is a canonical isomorphism between the states $|\phi\rangle \in \mathcal{H}_\gamma \otimes \overline{\mathcal{H}}_\gamma$ and $|\sigma\rangle \in \mathcal{H}_f \otimes \overline{\mathcal{H}}_f$ and operators $\phi : \mathcal{H}_\gamma \rightarrow \mathcal{H}_\gamma$ and $\sigma : \mathcal{H}_f \rightarrow \mathcal{H}_f$. Since $|\phi\rangle$ and $|\sigma\rangle$ are full-rank, these operators are invertible. We can therefore simply replace the central tensor in the subregion by $V'W|\phi\rangle|\sigma\rangle$, where

$$V' = \left(\prod_i \sigma_i^{-1} \phi_i^{-1} \right) V \tag{3.22}$$

and the product is taken over all RT surfaces on the boundary of a . In the case where there is a unique “uppermost” RT surface for our discretization, the isometry W will map directly to the boundary and will not require modification. In the case of a discretization where two or more RT surfaces both neighbor the root node, the isometry W will need to be modified on any of its outgoing legs that pass through already-constructed RT surfaces, as in equation (3.22).

By construction, the new state associated to the bulk subregion will approximate the old state, and since by assumption the rest of the tensor network forms an isometry from $V'W|\phi\rangle|\sigma\rangle$ (or $V'W'|\phi\rangle|\sigma\rangle$ in the case of multiple “uppermost” RT surfaces) to the global boundary, the new state constructed by the entire network will continue to approximate the target CFT state. The map V , which maps the new bulk subregion across the new RT surface into the rest of the network, is an exact isometry even though V' is not. This validates our assumption that extension maps that follow the root-leaf orientation of the network should always be isometries. By induction, we can therefore construct a tensor network state whose geometry agrees with an arbitrary tree discretization and which approximately reproduces the original CFT state.

The only claim that we have made but are yet to prove is that extension maps which flow partially against the root-leaf orientation of the underlying graph are still approximate, if not exact, isometries. This claim was not required as part of our inductive procedure for constructing tree networks, but will be useful for constructing more general networks in following sections. Consider an arbitrary RT surface in some tree network discretization. We have two extension maps, one in each direction, that map the edge state $|\phi\rangle|\sigma\rangle$ associated with this surface to the entire boundary. One map $V : \overline{\mathcal{H}}_\gamma \otimes \overline{\mathcal{H}}_f \hookrightarrow \mathcal{H}_{A^c}$ flows entirely with the orientation of the graph and so is an exact isometry. When the RT surface was first added to the network, the other extension map $W : \mathcal{H}_\gamma \otimes \mathcal{H}_f \hookrightarrow \mathcal{H}_A$ was also an exact isometry for exactly the same reasons. However, because this extension map flows partially against the root-leaf orientation of the graph, it will continue to change as additional surfaces are added to the network. This is in contrast to the “downwards” isometry V , which remains unaltered because of the order in which we chose to add the RT surfaces.

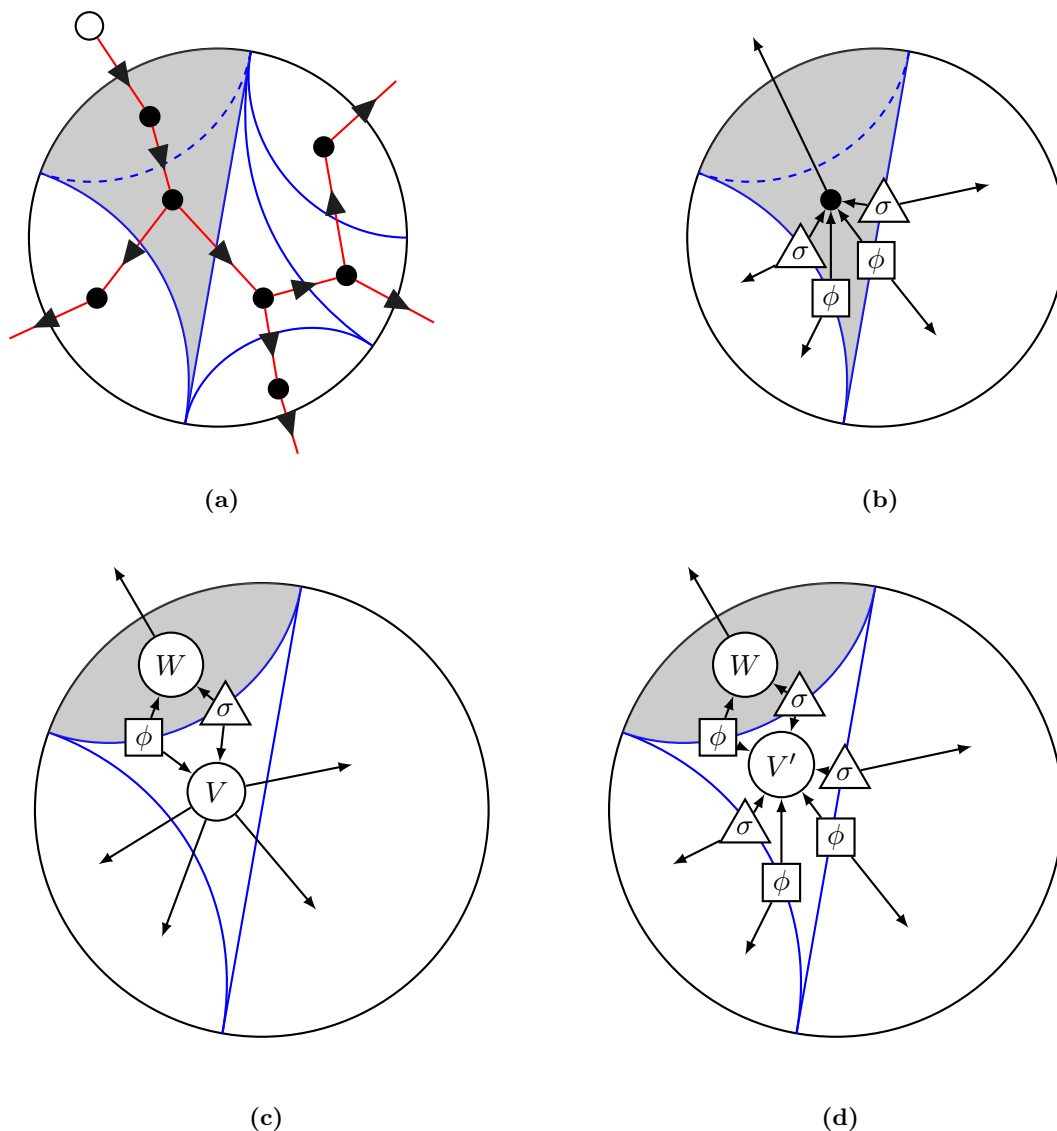


Figure 7. The final step of a tree tensor network construction for the bulk discretization of vacuum AdS_3 with orientation given in figure 5. (a) It is assumed that a tree network has been constructed for the discretization that consists of all but the final (dashed) RT surface. The bulk subregion a , shaded in gray, is the smallest bulk subregion containing this surface in its interior. (b) The state of the network on a is sketched explicitly. Since the final RT surface has not yet been added to the network, the bulk subregion a contains only one tensor. As all extension maps away from a follow the root-leaf orientation of the network, they are all exact isometries. (c) The subregion state on a is distilled across the new RT surface into a bipartite tensor network such as the one given in (3.15). (d) To preserve the structure of the states $|\phi_i\rangle$ and $|\sigma_i\rangle$ on the already-constructed RT surfaces, the bipartite isometry V is replaced by the map V' given in equation (3.22). Note that all extension maps that begin on the new smallest bulk subregion (shaded here) are still exact isometries.

In the final network, once all RT surfaces have been added, we call the extension map that flows partially against the root-leaf orientation $X : \mathcal{H}_\gamma \otimes \mathcal{H}_f \rightarrow \mathcal{H}_A$. This map will general not be an exact isometry; we will show that it is still an approximately isometry in an appropriate sense.

Because our construction is designed so that the tensor network approximately reproduces the boundary state at every stage in its construction, we find

$$V \otimes W|\phi\rangle|\sigma\rangle \approx V \otimes X|\phi\rangle|\sigma\rangle. \quad (3.23)$$

The left-hand side of (3.23) describes the state produced by the entire tensor network when the RT surface is first added to the network, while the right-hand side is the state of the final tensor network. Since $|\phi\rangle|\sigma\rangle$ is fully entangled (i.e., its reduced density matrices on either side of the RT surface are full-rank), exactness of (3.23) would imply that $X = W$ and so the extension map that flows partially against the network orientation would have to remain an exact isometry. As such, we can interpret (3.23) as showing that X is an approximate isometry with respect to a particular metric that is adapted to the state $|\phi\rangle|\sigma\rangle$. It can be equivalently written as

$$\|(X - W)\phi \otimes \sigma\|_2 = \|(X - W)\rho_\phi^{1/2}\rho_\sigma^{1/2}\|_2 \leq \varepsilon, \quad (3.24)$$

where $\|A\|_2 \equiv \sqrt{\text{Tr}(A^\dagger A)}$ is the Hilbert-Schmidt norm, $\rho_\phi\rho_\sigma = \phi^2\sigma^2$ is the reduced density matrix of $|\phi\rangle|\sigma\rangle$, and $\varepsilon > 0$ is small. This removes the isometry V from (3.23), and so gives a distance that depends only on X , W and the reduced density matrix $\rho_\phi\rho_\sigma$.

Finally, we can look at the partial trace of (3.23) over $\mathcal{H}_f \otimes \mathcal{H}_\gamma$. Using the fact that the trace norm $\|\rho\|_1 = \text{Tr} \sqrt{\rho^\dagger \rho}$ is monotonically decreasing under the partial trace, it follows from (3.23) that

$$\left\| \text{Tr}_{f\gamma} \left(X|\phi\rangle|\sigma\rangle\langle\phi|\langle\sigma|X^\dagger \right) - \rho_\phi\rho_\sigma \right\|_1 = \left\| \rho_\phi^{1/2}\rho_\sigma^{1/2}(X^\dagger X - \mathbf{1})\rho_\phi^{1/2}\rho_\sigma^{1/2} \right\|_1 \leq \varepsilon \quad (3.25)$$

This is a strictly weaker condition than (3.23) and (3.24): inefficiencies in the Fuchs-van de Graaf inequalities [31] mean that we cannot recover (3.23) and (3.24) without some loss of precision. However, it is perhaps the easiest condition to interpret.

Let $\{x_i\}$ be the eigenvalues of $X^\dagger X$. If $|\sigma\rangle$ were maximally mixed, we could rewrite (3.25) as

$$\frac{1}{d} \sum_i |x_i - 1| \leq \varepsilon. \quad (3.26)$$

In other words, we would understand (3.25) as saying that *on average* the eigenvalues of $X^\dagger X$ is close to one. Hence $X^\dagger X \simeq \mathbf{1}$ and X is an approximate isometry. Since $|\sigma\rangle$ is full rank but is not maximally mixed, we should instead think of (3.25) as a *weighted* average of $|x_i - 1|$. Of course, since in general $X^\dagger X$ and ρ_σ will not commute, this interpretation is not quite literal. It does, however, provide the correct intuition.

Thus far, we have only considered the approximate isometry condition for extension maps that go from an RT surface all the way to the boundary. What about an extension map X' that flows only partially through the network, whose image lies on RT surfaces in

addition to, or instead of, the boundary? In this case, the “output” RT surfaces may not have been added to the network when the “input” RT surface was added, so it may not necessarily be possible to compare the final extension map X' to some intermediary extension map W' as we did in equation (3.23). However, it will always be possible to *compose* X' with some exact isometry W' so that the resulting operator $W'X'$ is an extension map from the “input” RT surfaces of X' out to the global boundary.²¹

Since W' is an exact isometry, applying (3.25) to the map $W'X'$ yields the following inequality:

$$\left\| \rho_\phi^{1/2} \rho_\sigma^{1/2} (X'^\dagger X' - \mathbf{1}) \rho_\phi^{1/2} \rho_\sigma^{1/2} \right\|_1 = \left\| \rho_\phi^{1/2} \rho_\sigma^{1/2} (X'^\dagger W'^\dagger W' X' - \mathbf{1}) \rho_\phi^{1/2} \rho_\sigma^{1/2} \right\|_1 \leq \varepsilon \quad (3.27)$$

We conclude that extension maps that do not flow all the way to the boundary, such as X' , satisfy the same approximate isometry condition given in (3.25) for extension maps that *do* flow all the way to the boundary.

Unfortunately, the approximate isometry conditions (3.23), (3.24), (3.25), and (3.27) are not quite as powerful as one might want. They tell us that X and X' are close to isometries “on average”, but, because of the large dimensions of the Hilbert spaces involved, they do not say much about the “worst-case” error. For example, a single eigenvalue x_i of $X^\dagger X$ can be very far from one without making a large contribution to the averaged error (3.26).

More formally, to bound the “worst-case” error, we would want to bound the operator norm

$$\|X - W\|_\infty = \sup_{|\psi\rangle} \frac{\|(X - W)|\psi\rangle\|}{\|\psi\rangle\|}. \quad (3.28)$$

However, the tightest bound that we can place on this norm is

$$\|X - W\|_\infty \leq \|(X - W\phi \otimes \sigma)\|_\infty \|\phi^{-1} \otimes \sigma^{-1}\|_\infty \leq \varepsilon \|\phi^{-1} \otimes \sigma^{-1}\|_\infty, \quad (3.29)$$

where the first inequality follows from the submultiplicativity of the operator norm and the second follows from the monotonicity of the Schatten norms and (3.24). Since the operator norm of $\phi^{-1} \otimes \sigma^{-1}$ is quite large, naïvely satisfying

$$\|\phi^{-1} \otimes \sigma^{-1}\|_\infty = e^{O(S)}, \quad (3.30)$$

we cannot make ε small enough to make this a tight bound.

If the RT surfaces of a tree tensor network are added in an arbitrary order, rather than adding child surfaces before parent surfaces, then the potentially large error in (3.29) prevents us from guaranteeing that the final network correctly approximates the original boundary state. In this section, we avoided this possibility by imposing a root-leaf orientation on our construction; however, when constructing more complicated sub-AdS scale tensor networks in section 6, the issue of errors in our approximate isometries will arise once again without the guarantee of a root-leaf orientation to stop them from blowing up.

²¹In the final network, in general, the extension map W' from the output surfaces of X' to the global boundary will not be an exact isometry. However, it will always be possible to find some *intermediary* network, constructed as part of the inductive procedure, where W' is exact. Since both final and intermediary networks reproduce the boundary state to within tolerance ε , either can be used to prove equation (3.27).

However, in practice, (3.23) and (3.25) should ensure that “generic” small perturbations in the subregion states of our network that occur during an unoriented induction should only lead to a small error in the final state. Large errors would only occur if these perturbations were somehow finely tuned to blow up when mapped out to the boundary via extension maps. We therefore think it very likely that our distillation procedure will produce tensor networks that approximate the original holographic state correctly *regardless* of the order in which Ryu-Takayanagi surfaces are added to the network, and more generally will produce accurate sub-AdS scale tensor networks even when we no longer have the luxury of a root-leaf orientation to constrain accumulated errors precisely.

4 Bulk legs and holographic quantum error correction

Thus far, we have focused on constructing a tensor network for a single, arbitrary holographic CFT state. In the literature, however, it is common to consider tensor networks that describe not only a single holographic state, but an entire code subspace of holographic states. These tensor networks have some dangling bulk legs that turn the entire tensor network into a bulk-to-boundary map that encodes bulk excitations into a subspace of the boundary (see, e.g., [9] and [10]). Ideally, the map from the bulk to the boundary will have some appropriate error correcting properties that allow bulk operator reconstruction to be interpreted in the language of quantum error correction.

We discuss here, briefly, how our construction can be extended to create such code spaces. We shall focus on a code space that consists of a single bulk qubit, localized within a single bulk region a of the tree network discretization shown in figure 8a. The generalization to more complicated bulk code spaces in tree tensor networks is straightforward, and we would optimistically expect similar results to hold for the more general constructions of sections 5 and 6.

In much of the existing literature on holographic quantum error correcting codes (see, e.g., [9] and [10]), one begins with a particular tensor network construction and then shows that it has quantum error correcting properties analogous to those of AdS/CFT. As in the rest of this paper, we reverse this script. We begin with the known quantum error correcting properties of AdS/CFT [13], and especially entanglement wedge reconstruction [36–39], and show that they imply the existence of tensor network constructions for code spaces of *actual* holographic states, which behave in exactly the same way as the toy models that have previously appeared in the literature. We emphasize that this should not be taken as an independent proof of entanglement wedge reconstruction. Instead it demonstrates that the error-correction properties of AdS/CFT itself and of holographic tensor networks toy models are not merely analogous to one another, but are in fact different examples of exactly the same phenomenon.

Consider, for example, a 2-dimensional “code subspace” of states in AdS/CFT that can be built from a single starting state by applying low-energy unitary bulk operators in only a single bulk region a of some tree network discretization.²² Any Ryu-Takayanagi

²²If the energy of the bulk matter excitations is $O(1)$ in a G_N expansion, then the backreaction on the geometry is $O(G_N)$, resulting in an $O(1)$ change in the Bekenstein-Hawking entropy of the surfaces bounding the region a . This is smaller than the $O(G_N^{-1/2})$ fluctuations already present in our network, and can therefore be neglected without changing our error estimates on the bond dimensions of the network.

surface neighboring the bulk region a splits the boundary into two regions: region A , the entanglement wedge of which contains the bulk region a , and its complement A^c . This is sketched for a particular choice of bulk region and Ryu-Takayanagi surface in figure 8b. Because of the usual error correcting properties of AdS/CFT [13, 33, 36, 37], the reduced density matrix on A^c is approximately the same for every state in the code space. Moreover, any two orthogonal states in the code space will have reduced density matrices on A that are approximately orthogonal; the trace distance between them will be close to maximal. This is in fact equivalent to the fact that the reduced density matrix on A^c is the same for every state, including superpositions, in the code space (see, e.g., the weak decoupling duality described in [40, 41]).

Let us suppose we have successfully constructed a tree tensor network for a single state $|\psi\rangle$ in the code space using the techniques of section 3. Because the reduced density matrices on A^c are the same for every RT surface bounding a and every state in the code space, any state in the code subspace can be represented by a tensor network that is identical to the one constructed for $|\psi\rangle$ *except* that it differs in the tensor associated to the bulk subregion a . It follows by linearity that we can describe the entire two-dimensional code space adding a two-dimensional bulk leg to the tensor in region a (cf. figure 8c). The resulting tensor network can be interpreted as a map $T : \mathcal{H}_{\text{bulk}} \hookrightarrow \mathcal{H}_{\text{CFT}}$ from the bulk to the boundary. Furthermore, by using our freedom to choose an inner product on $\mathcal{H}_{\text{bulk}}$, we can ensure that T is an exact isometry.

Showing that our tensor networks are quantum error correcting in the sense of [9] requires showing that for any choice of RT surface bounding a and corresponding boundary subregion A , any bulk operator on $\mathcal{H}_{\text{bulk}}$ has an equivalent boundary operator whose support lies only on A . By “equivalent,” we mean that acting with one of these boundary operators on any state in the code subspace produces approximately the same state that one would obtain by acting with the original, bulk operator on $\mathcal{H}_{\text{bulk}}$. Such boundary representations of bulk operators are obtained in [9] by using the tensor network to “push the bulk operator through the network” and into the boundary.

To be precise, the map that will be used to push bulk operators to the boundary region A is the “controlled extension map” $X : \mathcal{H}_f \otimes \mathcal{H}_\gamma \otimes \mathcal{H}_{\text{bulk}} \rightarrow A$, sketched in figure 8c. We call this a *controlled extension map* because once a state $|\psi\rangle$ is fixed on the bulk leg $\mathcal{H}_{\text{bulk}}$, the resulting map $X|\psi\rangle$ is an extension map in the corresponding tensor network in the sense of section 3.2. We will show first that X is an approximate isometry in the sense of equation (3.25), and second that this condition allows us to use X to push bulk operators to equivalent operators on the boundary.

Let $|0\rangle, |1\rangle$ form a basis for $\mathcal{H}_{\text{bulk}}$. The controlled extension map X can then be represented in this basis as

$$X = X_0|0\rangle + X_1|1\rangle, \tag{4.1}$$

where X_0 and X_1 are the extension maps from the RT surface to A for the tree tensor networks defined by specifying bulk leg states $|0\rangle$ and $|1\rangle$ respectively. We showed in section 3.2 that extension maps in tree tensor networks are approximate isometries, and so

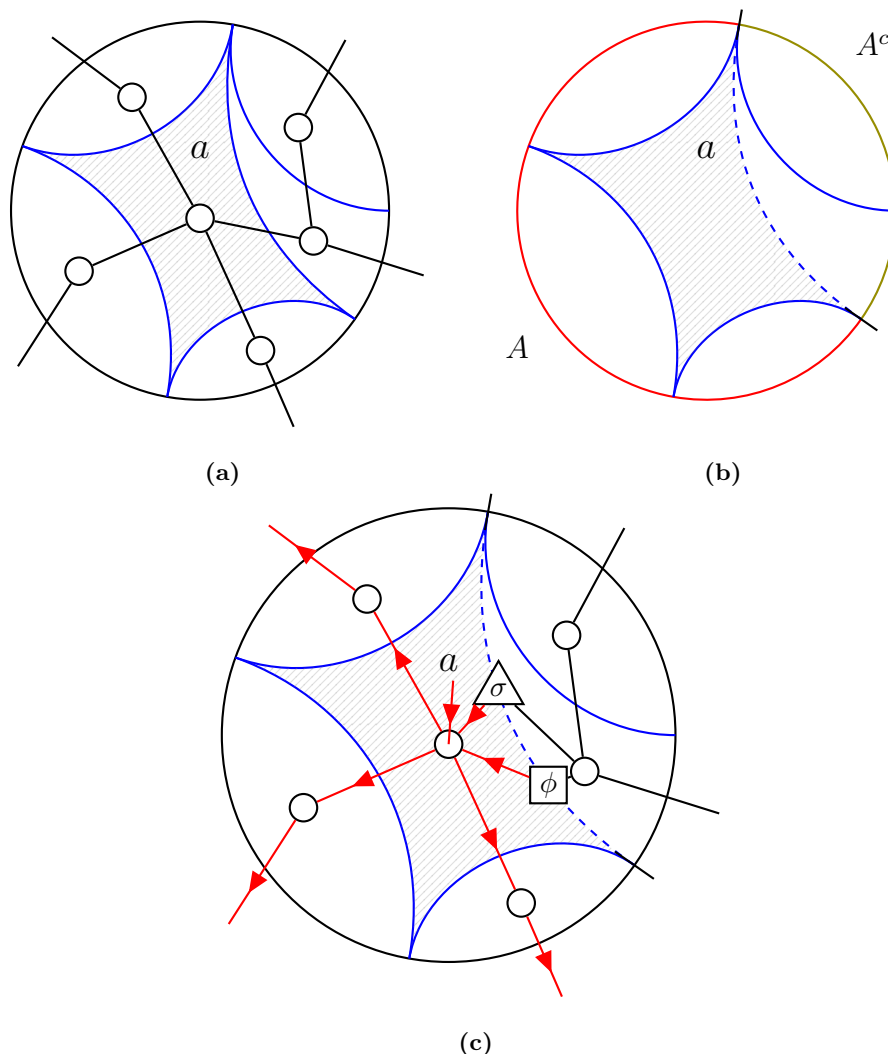


Figure 8. (a) A tree tensor network for a particular discretization of vacuum AdS_3 . A bulk subregion a of the discretization (shaded here) has been chosen to create a code subspace of excitations around the vacuum generated by operators whose support lies entirely within a . Note that the $|\phi\rangle$ and $|\sigma\rangle$ edge states have been suppressed in this sketch for the sake of visual clarity. (b) Choosing a particular RT surface that bounds a , denoted here with a dashed line, partitions the global boundary into two regions: A , which contains a in its entanglement wedge, and its complement A^c . (c) As explained in the text, a code subspace of excitations localized in a can be represented by adding a single bulk leg to the tensor in region a . For a particular choice of RT surface bounding a , the “controlled extension map”, which maps the bulk leg in a plus the edge states on the chosen RT surface into the boundary region A , is sketched with red arrows. In order to clarify where the extension map begins, the edge states $|\phi\rangle$ and $|\sigma\rangle$ are shown explicitly on the chosen RT surface, and suppressed on all other RT surfaces in the diagram.

X_0 and X_1 satisfy

$$X_0^\dagger X_0 \simeq X_1^\dagger X_1 \simeq \mathbb{1} \tag{4.2}$$

in the sense of equation (3.25). Because orthogonal bulk states are almost orthogonal in region A , these extension maps satisfy

$$X_1^\dagger X_0 \simeq X_0^\dagger X_1 \simeq 0 \tag{4.3}$$

in the same sense. It follows immediately that $X^\dagger X \simeq \mathbb{1}$ and so X is an approximate isometry in the sense of (3.25). By the same arguments as in section 3.2, we can also show that controlled extension maps that end on RT surfaces are also approximate isometries.

Given an operator \hat{O}_b acting on $\mathcal{H}_{\text{bulk}}$, we wish to use the fact that X is an approximate isometry to produce a boundary operator \hat{O}_A supported on A whose action on the code subspace is approximately the same as \hat{O}_b . More precisely, for a state $|\psi\rangle \in \mathcal{H}_{\text{bulk}}$, we wish to show

$$\hat{O}_A T|\psi\rangle \approx T\hat{O}_b|\psi\rangle. \tag{4.4}$$

This is exactly the same sense of bulk reconstruction through quantum error correction that was developed for exact tensor network toy models in [9]. In terms of the controlled extension map X and the bulk tensor network state $|T_{A^c}\rangle$ associated to the entanglement wedge of A^c , the tensor network map T can be decomposed as

$$T|\psi\rangle = X|T_{A^c}\rangle|\psi\rangle. \tag{4.5}$$

Equation (4.4) therefore becomes

$$\hat{O}_A X|T_{A^c}\rangle|\psi\rangle \approx X\hat{O}_b|T_{A^c}\rangle|\psi\rangle. \tag{4.6}$$

As in [9], we can find a boundary representation of a bulk operator by simply conjugating with X , i.e.,

$$\hat{O}_A \equiv X\hat{O}_bX^\dagger. \tag{4.7}$$

Using the approximate isometry condition that $X^\dagger X$ acts approximately as the identity on $|T_{A^c}\rangle|\psi\rangle$ for any bulk state $|\psi\rangle$, we see that \hat{O}_A satisfies equation (4.4). If X were an exact isometry, as in [9], the map from bulk operators to boundary operators given in (4.7) would be unital on the image of X (i.e., it maps the identity on $\mathcal{H}_{\text{bulk}}$ to the boundary projector onto the image of X). This condition is desirable from the perspective of quantum information theory, as it ensures that the map given in (4.7) is a quantum channel in the Heisenberg picture.

In order to ensure that this condition holds for our approximate isometries, thus making the bulk-to-boundary operator map into a genuine quantum channel, we instead define the boundary representation of a bulk operator as

$$\hat{O}_A \equiv X(X^\dagger X)^{-1/2}\hat{O}_b(X^\dagger X)^{-1/2}X^\dagger. \tag{4.8}$$

From the isometry condition (3.25), we see that \hat{O}_A still satisfies equation (4.4), and thus this an approximate boundary representation of the bulk operator \hat{O}_b . We can also see

plainly that the bulk-to-boundary operator map given by equation (4.8) maps the identity on $\mathcal{H}_{\text{bulk}}$ to the identity on the image of X , and is thus a quantum channel on operators in the code subspace (i.e., a completely positive, unital map on operators). We conclude that every bulk operator \hat{O}_b has an equivalent boundary operator representation \hat{O}_A whose support lies only on A and whose action on the code subspace is approximately the same as the action of \hat{O}_b , as desired.

The boundary operator obtained via the quantum channel given in equation (4.8) may be very different from the boundary operator on the same region obtained through the extrapolate dictionary and the HKLL reconstruction procedure [42]; however, these operators will agree on the boundary code subspace obtained by taking the image of $\mathcal{H}_{\text{bulk}}$ under the map defined by the tensor network. Our tensor network description of a state in the AdS/CFT correspondence therefore displays exactly the same bulk-to-boundary operator mapping properties of the celebrated HaPPY class of holographic codes [9]. We conclude that tensor networks models of bulk operator reconstruction may be taken quite literally. They are not just toy models!

5 Loop tensor networks from holographic entanglement of purification

The tree networks constructed in section 3 already constitute a large class of tensor networks for the AdS/CFT correspondence, but we might still hope for a more general construction. A single bond in a tree tensor network corresponds to a complete Ryu-Takayanagi surface in the bulk spacetime from which the network was constructed, so the information contained in that bond is *a priori* distributed nonlocally across the entire corresponding surface. In order to localize degrees of freedom at sub-AdS scales within the bulk geometry, we must therefore find a way to divide each bond geometrically along its corresponding Ryu-Takayanagi surface. In the bulk discretization picture of section 3, this corresponds to constructing a holographic tensor network with *loops*.

To construct such a tensor network, we will need to understand the *holographic entanglement of purification* [15, 16], a geometric quantity in the bulk spacetime that is conjectured to correspond to an information-theoretic quantity involving a tensor factorization of the information on an RT surface. In this section, we will assume the holographic entanglement of purification conjecture, and then use it to construct a geometrically accurate tensor network with *a single loop* for an arbitrary holographic CFT state. With some further assumptions introduced in section 6, we will then be able to extend this procedure to construct tensor networks with arbitrarily many loops, and therefore to localize information in the bulk arbitrarily well within the regime of validity of the Ryu-Takayanagi formula (i.e., above the string/Planck length scales).

5.1 Holographic entanglement of purification

For a quantum state $\rho_{(AB)}$ on a bipartite Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$, the entanglement of purification [43] between subsystems A and B is defined as

$$E_P(A : B) = \inf_{|\Psi\rangle_{AA'BB'}} S(AA'), \tag{5.1}$$

where $|\Psi\rangle_{AA'BB'} \in \mathcal{H}_A \otimes \mathcal{H}_{A'} \otimes \mathcal{H}_B \otimes \mathcal{H}_{B'}$ is a purification of $\rho_{(AB)}$. The infimum in equation (5.1) is taken over all possible purifications of $\rho_{(AB)}$ in all possible auxiliary Hilbert spaces $\mathcal{H}_{A'}$ and $\mathcal{H}_{B'}$.

If $\rho_{(AB)}$ is a mixed state, then the von Neumann entropy $S(A)$ no longer measures the entanglement (or even the correlation) between A and B in any meaningful sense, since some portion of $S(A)$ is inherited from the nonzero von Neumann entropy of $\rho_{(AB)}$. The von Neumann entropy $S(A)$ may be non-zero for a product state $\rho_{(A)} \otimes \rho_{(B)}$, which has no correlation between the two subsystems. The entanglement of purification E_P is a somewhat better measure of the degree to which A and B are entangled (or at least correlated), as it measures the minimal entanglement between AA' and BB' for any purification. As a result, the entanglement of purification is zero for product states, and is non-increasing under local operations. It is not a true entanglement monotone in the sense of [44], however, since it may be non-zero even for separable states (which only have classical correlation), and it may be increased by classical communication.

As with other information-theoretic quantities, one might hope that the entanglement of purification has a geometric dual in the context of AdS/CFT. For two subregions A_1 and A_2 of a holographic CFT state, it has recently been conjectured in [15, 16] that $E_P(A_1 : A_2)$ is given to leading order in G_N by the area of the *entanglement wedge cross-section*, $E_W(A_1 : A_2)$. Formally, $E_W(A_1 : A_2)$ is defined as

$$E_W(A_1 : A_2) = \frac{\text{area}(\Sigma_{A_1:A_2})}{4G_N}, \tag{5.2}$$

where $\Sigma_{A_1:A_2}$ is the minimal surface anchored to the Ryu-Takayanagi surface of $A_1 \cup A_2$ that partitions the entanglement wedge into a portion whose boundary contains all of A_1 and a disjoint portion whose boundary contains all of A_2 . This surface is sketched in figure 9 in vacuum AdS_3 for two different typical configurations of boundary regions A_1 and A_2 .

If A_1 and A_2 are connected subregions of some larger connected region $A = A_1 \cup A_2$ (i.e., if A_1 and A_2 form a connected partition of A as in figure 9a, then we call the state $|\Psi\rangle_{A_1 A'_1 A_2 A'_2}$ that saturates the infimum in equation (5.1) the *minimally entangled purification* (MEP) of the partition $A_1 : A_2$.^{23,24}

The holographic entanglement of purification conjecture has been generalized to multipartite and conditional entropies (and their geometric duals) in [46–49]. For the moment, however, we will concern ourselves only with the minimally entangled purification of a connected partition $A_1 : A_2$, and the associated entanglement of purification $E_P(A_1 : A_2)$.

²³Of course, since (5.1) contains an infimum rather than a minimum, it is not necessarily saturated by any state $|\Psi\rangle_{A_1 A'_1 A_2 A'_2}$. In this case, we take the MEP to be a fixed state that saturates the infimum to within tolerance ϵ . Since one might expect that the infimum in (5.1) could be saturated arbitrarily well by taking the purifying spaces A'_1 and A'_2 to be arbitrarily large, allowing this finite tolerance in the MEP prevents the dimensions of A'_1 and A'_2 from blowing up.

²⁴A similar object, defined by a procedure in which one is also permitted to minimize over all possible partitions $A_1 : A_2$, has previously been studied in [45].

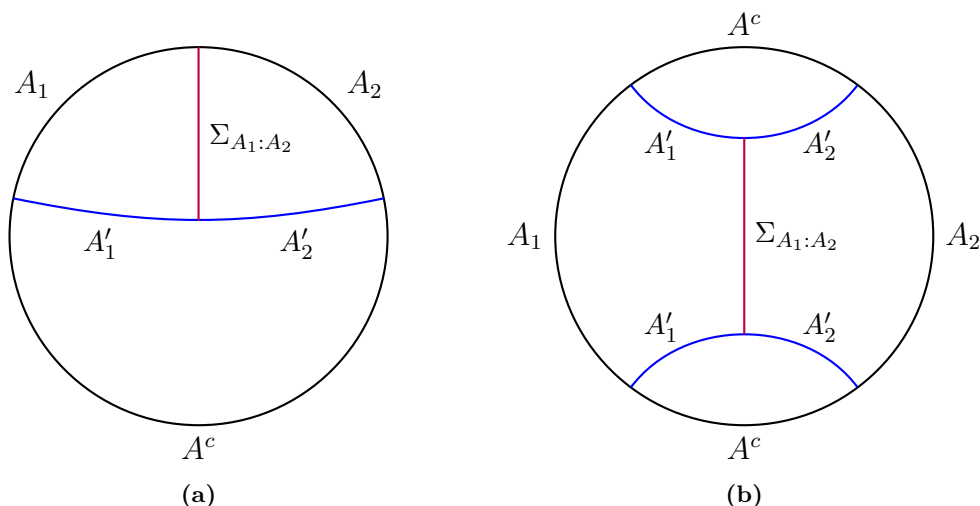


Figure 9. The entanglement wedge cross-section for two different configurations of boundary subregions A_1 and A_2 in vacuum AdS_3 . In both (a) and (b), the entanglement wedge cross-section divides the entanglement wedge (the bulk region bounded by A_1 , A_2 , and RT surface of $A_1 \cup A_2$) into two disjoint regions, each of which contains only one of either A_1 or A_2 in its boundary. The subregions of each RT surface inherited from this partition have been instructively labeled A'_1 and A'_2 for reasons explained in section 5.2.

5.2 Tensor networks from minimally entangled purifications

Consider a holographic CFT state $|\psi\rangle$ with some subregion A that is further divided into a connected partition $A_1 : A_2$, as sketched in figure 9a. The reduced state $\rho_{(A_1 A_2)}$ has a minimally entangled purification $|\Psi\rangle_{A_1 A'_1 A_2 A'_2}$ chosen so that $S(A_1 A'_1) = E_P(A_1 : A_2)$. According to the holographic entanglement of purification conjecture, this implies

$$S(A_1 A'_1) = E_W(A_1 : A_2) = \frac{\text{area}(\Sigma_{A_1:A_2})}{4G_N}. \quad (5.3)$$

Examining figure 9a, it is easy to see that $\Sigma_{A_1:A_2}$ is the minimal bulk surface anchored to the boundary region A_1 and a subregion of the Ryu-Takayanagi surface of $A_1 \cup A_2$, which we have instructively labeled A'_1 . Equation (5.3) looks just like the Ryu-Takayanagi formula (2.8) if one supposes that $|\Psi\rangle_{A_1 A'_1 A_2 A'_2}$ is a holographic state in some boundary theory with a domain corresponding to the codimension-2 bulk surface made up of A_1 and A_2 along with the Ryu-Takayanagi surface of their union.

Indeed, as explained in [15, 16], it makes sense to interpret the MEP $|\Psi\rangle_{A_1 A'_1 A_2 A'_2}$ as a geometric state on some subregion of the bulk, where the auxiliary Hilbert spaces $\mathcal{H}_{A'_1}$ and $\mathcal{H}_{A'_2}$ are identified with the bulk surfaces A'_1 and A'_2 , respectively. The conjectured surface-state correspondence [50] of holographic CFTs suggests that any codimension-2 surface in a holographic spacetime corresponds to some state in a boundary theory that encodes the physics in the entanglement wedge of that surface.²⁵ If this is to be believed, then

²⁵The “entanglement wedge” of a codimension-2 bulk surface Σ is defined in analogy with the entanglement wedge of a boundary interval as the subregion of the bulk bounded by Σ and by the minimal surface homologous to Σ that shares its boundary.

it is natural to assume that the MEP of the partition $A_1 : A_2$ corresponds to a state on the codimension-2 bulk surface $\Sigma_{A_1 A'_1 A_2 A'_2} = A_1 \cup A_2 \cup A'_1 \cup A'_2$.

To make this proposal precise, we assume that the entanglement entropies of subsystems of the MEP are given at leading order by the Ryu-Takayanagi formula applied to the codimension-2 surface $\Sigma_{A_1 A'_1 A_2 A'_2}$. In particular, since A'_1 and A'_2 are subregions of a minimal surface and hence minimal themselves, this assumption tells us that A'_1 and A'_2 are themselves the Ryu-Takayanagi surfaces for the Hilbert Space factors $\mathcal{H}_{A'_1}$ and $\mathcal{H}_{A'_2}$. In other words, their von Neumann entropies satisfy:

$$S(A'_1) = \frac{\text{area}(A'_1)}{4G_N} + o\left(\frac{1}{G_N}\right), \quad (5.4)$$

$$S(A'_2) = \frac{\text{area}(A'_2)}{4G_N} + o\left(\frac{1}{G_N}\right). \quad (5.5)$$

We require one more assumption to build a tensor network using the holographic entanglement of purification, which is that the MEP has the same smooth min- and max-entropy properties (2.30) and (2.31) as holographic CFT states. In section 2.3, we argued that this will be true for any state with von Neumann entropies given by equations of the form (5.4) and (5.5), and with extensive Rényi entropies that take the geometrical form given in equation (2.16). If the MEP is indeed a holographic state for the portion of the bulk bounded by $\Sigma_{A_1 A'_1 A_2 A'_2}$, then it should satisfy both of these properties.

If the MEP has entropies given to leading order by the Ryu-Takayanagi formula on subregions of $\Sigma_{A_1 A'_1 A_2 A'_2}$, and if the smooth min- and max-entropies of those subregions are also given by (2.30) and (2.31), then the MEP satisfies all the conditions required to build a tree tensor network (cf. section 3).²⁶ This means that we can build a tree network for the MEP that matches the bulk geometry contained within $\Sigma_{A_1 A'_1 A_2 A'_2}$. Specifically, there exists some state $|\Psi_{(\varepsilon)}\rangle_{A_1 A'_1 A_2 A'_2}$ that approximates the MEP with high fidelity with a tensor network representation

$$\begin{aligned} \Psi_{(\varepsilon)}^{A_1 A'_1 A_2 A'_2} &= T^{A'_1}_{\gamma_1 f_1} U^{A_1}_{\bar{\gamma}_1 \bar{f}_1 \gamma_2 f_2} V^{A_2}_{\bar{\gamma}_2 \bar{f}_2 \gamma_3 f_3} W^{A'_2}_{\bar{\gamma}_3 \bar{f}_3} \\ &\quad \times \phi_{(1)}^{\gamma_1 \bar{\gamma}_1} \phi_{(2)}^{\gamma_2 \bar{\gamma}_2} \phi_{(3)}^{\gamma_3 \bar{\gamma}_3} \sigma_{(1)}^{f_1 \bar{f}_1} \sigma_{(2)}^{f_2 \bar{f}_2} \sigma_{(3)}^{f_3 \bar{f}_3} \end{aligned} \quad (5.6)$$

with bond dimensions given by:

$$\dim \mathcal{H}_{\gamma_1} = e^{S(A'_1) - O(\sqrt{S(A'_1)})}, \quad (5.7)$$

$$\dim \mathcal{H}_{f_1} = e^{O(\sqrt{S(A'_1)})}, \quad (5.8)$$

$$\dim \mathcal{H}_{\gamma_2} = e^{S(A_1 A'_1) - O(\sqrt{S(A_1 A'_1)})}, \quad (5.9)$$

$$\dim \mathcal{H}_{f_2} = e^{O(\sqrt{S(A_1 A'_1)})}, \quad (5.10)$$

$$\dim \mathcal{H}_{\gamma_3} = e^{S(A'_2) - O(\sqrt{S(A'_2)})}, \quad (5.11)$$

$$\dim \mathcal{H}_{f_3} = e^{O(\sqrt{S(A'_2)})}. \quad (5.12)$$

²⁶As in section 3.2, one must choose a root-leaf orientation for the MEP network in order to determine the order in which its subregions should be distilled. The choice of ordering will not matter for our purposes, though as usual it will determine which isometries are exact and which are only approximate.

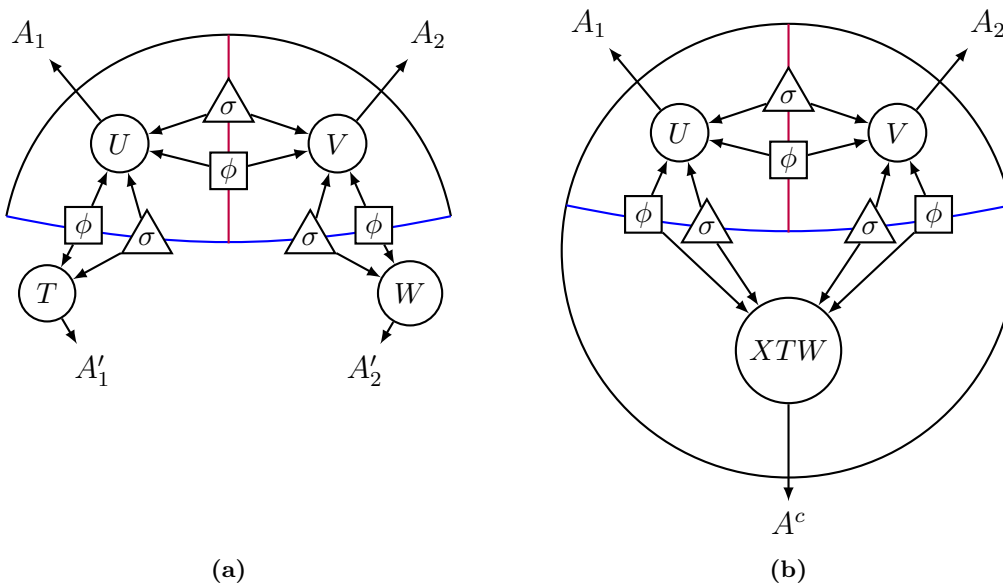


Figure 10. (a) A tensor network for the minimally entangled purification of two neighboring boundary regions in vacuum AdS_3 , as given in equation (5.6). (b) The one-loop tensor network for the full boundary state given by equation (5.16), obtained by embedding the MEP isometrically into the global boundary. Note that in both figures, subscripts on the $\phi_{(i)}$ and $\sigma_{(i)}$ edge states have been suppressed.

The outer product expression given in (5.6) is not particularly illuminating on its own, but has a natural interpretation in the geometric picture of tensor networks. In figure 10a, this tensor network is shown superposed over the geometric picture of the minimally entangled purification, with each network bond passing through its corresponding bulk surface. Equations (5.3), (5.4), and (5.5) imply that the bond dimensions of this tensor network match the areas of the surfaces A'_1 , A'_2 , and $\Sigma_{A_1:A_2}$ in figure 9a, justifying our interpretation of expression (5.6) as an approximate tensor network for the minimally entangled purification that matches the geometric properties of its holographic dual.

The tree network for the minimally entangled purification constitutes a tensor network for the “top half” of the bulk discretization shown in figures 9a and figure 10a. To find a geometric tensor network for the full boundary state, we need to find a tensor corresponding to the bulk region that lies between $A'_1 \cup A'_2$ and the complementary boundary region A^c . Since the MEP $|\Psi\rangle_{A_1 A'_1 A_2 A'_2}$ and the original global boundary state $|\psi\rangle_{A_1 A_2 A^c}$ are both purifications of the reduced boundary state $\rho_{(A_1 A_2)}$, they are related by an isometry on the purifying space. That is, there exists an isometry²⁷

$$X : \mathcal{H}_{A'_1} \otimes \mathcal{H}_{A'_2} \hookrightarrow \mathcal{H}_{A^c} \quad (5.13)$$

such that

$$|\psi\rangle_{A_1 A_2 A^c} = (I_{A_1 A_2} \otimes X) |\Psi\rangle_{A_1 A'_1 A_2 A'_2}. \quad (5.14)$$

²⁷Formally, this isometry only exists if $\dim(\mathcal{H}_{A^c}) \geq \dim(\mathcal{H}_{A'_1} \otimes \mathcal{H}_{A'_2})$. Since the dimension of the latter space is given to leading order by $e^{S(A^c)}$, and $S(A^c)$ is much smaller than $\log \dim(\mathcal{H}_{A^c})$, this bound is satisfied here and the isometry exists.

In tensor notation, this is

$$\psi^{A_1 A_2 A^c} = X^{A^c}_{A'_1 A'_2} \Psi^{A_1 A'_1 A_2 A'_2}. \quad (5.15)$$

Since X is an isometry, the minimally entangled purification $|\Psi\rangle_{A_1 A'_1 A_2 A'_2}$ can be replaced with the nearby state $|\Psi_{(\varepsilon)}\rangle_{A_1 A'_1 A_2 A'_2}$ without any additional loss of precision. That is, the state

$$\psi_{(\varepsilon)}^{A_1 A_2 A^c} = X^{A^c}_{A'_1 A'_2} \Psi_{(\varepsilon)}^{A_1 A'_1 A_2 A'_2}. \quad (5.16)$$

approximates ψ as well as $\Psi_{(\varepsilon)}$ approximates Ψ . Plugging the tensor network expression for $\Psi_{(\varepsilon)}$ from equation (5.6) into equation (5.16) (and contracting the tensors X , T , and W into a single bulk tensor XTW) yields a tensor network description for the full (approximate) boundary state $\psi_{(\varepsilon)}$ with bond dimensions matching the areas of the surfaces A'_1 , A'_2 , and $\Sigma_{A_1:A_2}$. This complete tensor network is sketched in figure 10b. By using the natural properties of the minimally entangled purification arising from the holographic entanglement of purification conjecture, we have managed to localize degrees of freedom within the Ryu-Takayanagi surface and hence to construct a non-tree tensor network for a generic holographic state with extremely high fidelity.

With this tensor network now fully constructed, we must now ask an important question about its extension maps: in what sense can they be shown to be isometries? Are they exact isometries, as in section 3.1? Are they merely approximate isometries in the sense of section 3.2? Or are they in fact neither of these things? Because the answer to this question will prove important both in constructing sub-AdS scale networks in section 6 and in formulating the no-go theorem that we prove in section 7.3, we shall answer this question systematically for each extension map of the network shown in figure 10b.

Most of the extension maps in figure 10b (e.g. the maps outwards from $\Sigma_{A_1:A_2}$ to $A_1 \cup A'_1$ and $A_2 \cup A'_2$ and the maps upwards from A'_1 to $A_1 \cup \Sigma_{A_1:A_2}$ and from A'_2 to $A_1 \cup \Sigma_{A_1:A_2}$) were also extension maps in the tree tensor network for the minimally entangled purification. As such, they are all at least approximate isometries in the sense of section 3.2; depending on the order in which RT surfaces were added to the tree network for the MEP, several of them will be exact isometries. The extension map XTW flowing downwards from the horizontal RT surface $A'_1 \cup A'_2$ is not an extension map from the tree tensor network for the MEP; however, since it is a composition of an exact isometry X with two MEP extension maps T and W , and since T and W are exact isometries regardless of the order in which the edges in the tree tensor network for the MEP were distilled, XTW is still an exact isometry in the final network.

There is one remaining extension map that one might also hope would be an isometry: the map upwards from $A'_1 \cup A'_2$ to $A_1 \cup A_2$. In this case, we have no good argument that it should be an exact, or even an approximate, isometry. In particular, the $|\phi\rangle$ and $|\sigma\rangle$ edge states on A'_1 and A'_2 were constructed to approximate the reduced density matrices of the MEP on $\mathcal{H}_{A'_1}$ and $\mathcal{H}_{A'_2}$ *individually*. The upwards extension map will be an approximate isometry if (and more importantly *only* if) the product of these reduced density matrices approximates the reduced density matrix of the entire MEP on $\mathcal{H}_{A'_1} \otimes \mathcal{H}_{A'_2}$. Unfortunately, the assumptions we have made about the MEP thus far are only sufficient to show that the mutual information $I(A'_1 : A'_2)$ is subleading in G_N ; they do not imply that it is

zero. We have no solid reason to believe that the reduced density matrix on $\mathcal{H}_{A'_1} \otimes \mathcal{H}_{A'_2}$ is actually close to a product state with respect to the trace norm. Furthermore, we will see in section 7.3 that there is good reason to think that this *cannot* be the case.²⁸

Henceforth, we shall refer to an extension map of this kind as a “moral” isometry. In using this terminology, we mean that — even though the extension map from $A'_1 \cup A'_2$ to $A_1 \cup A_2$ in figure 10b is unlikely to be an exact or even approximate isometry — we expect that small errors introduced in the “bottom-half” tensor XTW will not blow up dramatically when mapped through the moral isometry to alter the “top-half” boundary state on $A_1 \cup A_2$. Since the moral isometry preserves the normalization of the full-rank state $|\phi_{(1)}\rangle|\sigma_{(1)}\rangle|\phi_{(2)}\rangle|\sigma_{(2)}\rangle$, and also preserves the entanglement entropy of this state to leading order, it is tempting to think of the moral isometry as a combination of an exact isometry and some other, non-isometric operator that acts only on a subleading number of degrees of freedom. We will revisit the issue of moral isometries in section 6, where we invoke the moral isometry condition to argue that distilling the bottom-half tensor XTW into a tree tensor network of its own will result in a global tensor network that still approximately reproduces the original boundary state of the CFT.

5.3 Multiple loops

Of course, nothing in our construction thus far has limited us to considering the case of a bipartite boundary partition $A_1 : A_2$. One might equally well wish to consider a more general multipartite partition, where a boundary region A is partitioned into n connected subregions as $A_1 : A_2 : \dots : A_n$. For simplicity, we assume that each subregion A_i only shares a boundary with at most two neighbors, A_{i-1} and A_{i+1} . In $2 + 1$ spacetime dimensions, any connected partition can be ordered such that this is true. By analogy with the holographic entanglement of purification conjecture, one would expect that minimizing the quantity

$$S(A_1 A'_1) + S(A_1 A'_1 A_2 A'_2) + \dots + S(A_1 A'_1 \dots A_{n-1} A'_{n-1}) \tag{5.17}$$

over all possible purifications into spaces $A'_1 \otimes \dots \otimes A'_n$ would correspond to minimizing the areas of surfaces that partition the entanglement wedge of $A = A_1 \cup \dots \cup A_n$ into n distinct subregions.²⁹ Each term in (5.17) should correspond to the area of one entanglement wedge cross-section, and since minimal surfaces cannot cross one another, the minimization of each area in the bulk can be performed independently up to subleading corrections. If our extended conjecture is correct, this implies that each term in the entropy sum (5.17) can also be minimized independently, at leading order.

²⁸The fact that this map is not even an approximate isometry is somewhat problematic for interpreting the bottom-half state of the tensor network in figure 10b in terms of the surface-state correspondence. It implies that the bottom half state is *not* an approximate purification of the reduced density matrix of the original holographic state on \mathcal{H}_{A^c} . However, we still hope that the relevant smooth entropies will behave correctly.

²⁹In [47], a slightly different notion of multipartite entanglement of purification was shown to satisfy constraining inequalities involving linear combinations of entanglement entropies, which are also satisfied by the corresponding geometric quantity. Similar proof techniques, both holographic and information theoretic, should suffice to place analogous constraints on the quantity given in equation (5.17). Since the holographic entanglement of purification conjecture was originally motivated in [15] by showing that the geometric bulk quantities and the information-theoretic boundary quantities satisfy the same constraints, our construction is equally well-motivated.

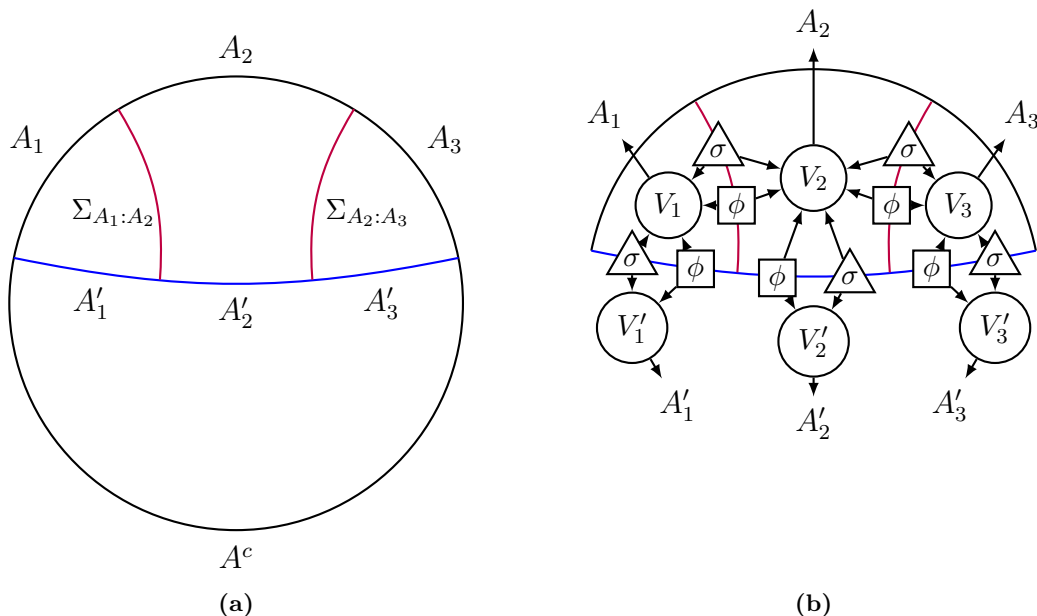


Figure 11. A rough sketch of the procedure for constructing a “3-to-one” network for a holographic CFT state. The boundary subregion A is given a connected partition $A_1 : A_2 : A_3$, and a minimally entangled purification for this partition is found by minimizing the sum $S(A_1 A'_1) + S(A_1 A'_1 A_2 A'_2)$. In (a), this partition is shown along with the surfaces $\Sigma_{A_1:A_2}$ and $\Sigma_{A_2:A_3}$ whose minimal areas should correspond to the minimization of this sum. If the MEP has entropies corresponding to areas of surfaces in the entanglement wedge of $A_1 \cup A_2 \cup A_3$, then a “top-half” tree tensor network can be constructed as in (b), which can then be completed by an isometry on the purifying space to obtain a tensor network for the global CFT state.

In the special case of two subregions (i.e., $n = 2$), the quantity given in (5.17) reduces to the one minimized in defining the entanglement of purification (5.1). If the MEP of this partition has entropies given by the areas of extremal surfaces contained within the entanglement wedge, then one can readily construct an “ n -to-one” network for the CFT state $|\psi\rangle$ by repeating the procedure detailed above. This is sketched for $n = 3$ in figure 11.

Note that for $n \geq 3$, the n -to-one network contains bonds that correspond to extremal surface subregions with areas that remain finite even when the ultraviolet CFT regulator and corresponding bulk radius regulator are removed. If the holographic entanglement of purification conjecture and its extensions hold down to the string and Planck scales, as assumed, then the areas of these subregions can be chosen to be arbitrarily small relative to the AdS scale by choosing suitably small boundary subregions (so long as we remain above the string and Planck scales).³⁰ The tensor network would therefore be capturing the bulk geometry at sub-AdS scales.

³⁰One important exception occurs when two neighboring entanglement wedge cross-sections $\Sigma_{A_1:A_2}$ and $\Sigma_{A_2:A_3}$ undergo a “phase transition” in the sense that they “jump” discontinuously across the RT surface even when the middle boundary region A_2 is made arbitrarily small. In this case, the region of the RT surface that is “skipped over” by this phase transition cannot be directly probed by a single application of the holographic entanglement of purification, and may in fact have an area well above the AdS scale. Nevertheless, techniques in section 6 should still allow us to construct a sub-AdS tensor network within such a region by using multiple, iterated applications of the holographic entanglement of purification conjecture. (This complication does not arise for bulk geometries that are close enough to a 2+1 AdS vacuum.)

Thus far, we have implicitly assumed that we are working in 2+1 spacetime dimensions. While the construction detailed above will certainly work in higher-dimensional spacetimes, it would no longer be completely accurate to claim that the n -to-one network captures the geometry at sub-AdS scales; while RT surface subregions can be chosen with finite, sub-AdS width, they also have at least one transverse direction that extends all the way to the boundary of the spacetime. Localizing the information on a single Ryu-Takayanagi surface to bounded, sub-AdS bulk regions in higher-dimensional spacetimes is a subtle procedure, and requires techniques from section 6. We will therefore comment on this generalization briefly in section 6.2.

6 Iteration and sub-AdS locality

In section 3, we showed that the Ryu-Takayanagi formula, together with constraints on the smooth min- and max-entropies that follow from the extensive growth of the ordinary Rényi entropies, is sufficient to construct geometrically appropriate tensor networks corresponding to an arbitrary discretization of the bulk by non-intersecting extremal surfaces. The resulting tensor networks are always tree tensor networks, where each bond of the network is associated to an entire extremal surface. In tree tensor networks, information is never localized within a single Ryu-Takayanagi surface.

In section 5, however, we were able to show that the holographic entanglement of purification conjecture can be used to associate network bonds to subregions of a single Ryu-Takayanagi surface. By assuming a natural extension of the holographic entanglement of purification conjecture to multipartite partitions of the boundary, these subregions could be made to have finite size even when the CFT and bulk regulators are removed. If the holographic entanglement of purification conjecture for multipartite boundary partitions holds up to stringy and quantum corrections, then these extremal surface subregions can be made arbitrarily small compared to the AdS scale in the semiclassical limit $G_N \rightarrow 0$ and $\lambda \rightarrow \infty$.³¹

We would like to go further by achieving some form of sub-AdS locality not only in the sense of dividing bonds along a single Ryu-Takayanagi surface, but in the general granularity of the network. More precisely, we would like to construct tensor networks where each tensor is associated to a bulk subregion that occupies a volume well below ℓ_{AdS}^{d-1} . We approach this problem by proposing a procedure to construct a tensor network for discretizations of the bulk whose discretization scale lies well below ℓ_{AdS} .

6.1 The four-tensor network

We begin by considering the simplest tensor network that our prior techniques were unable to address, namely the four-tensor square network shown in figure 12a that corresponds to a discretization of the bulk by two complete, intersecting extremal surfaces. One way to construct such a network involves a process of iteration: begin by constructing a one-loop network like the one shown in figure 10b, where the “bottom half” of the discretization in 12a is represented by a single tensor, then divide this tensor into two tensors that

³¹Note that since the bond dimension of the corresponding network edge goes like $e^{\text{area}/4G_N}$, the bond dimension will still diverge in the semiclassical limit.

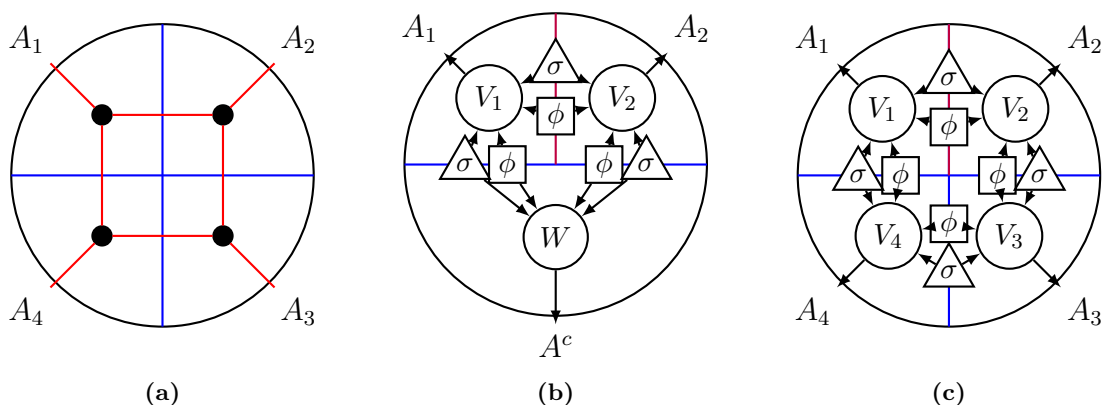


Figure 12. Constructing a tensor network for the four-tensor discretization by an iterative procedure. (a) A discretization of the bulk into four regions by two intersecting Ryu-Takayanagi surfaces, along with the corresponding dual graph. (b) The one-loop network for the boundary partition $A_1 : A_2$ obtained from the minimally entangled purification of this partition. (c) The full four-tensor network, obtained from (b) by distilling entanglement out of the bottom tensor W .

represent the bulk subregions in the discretization. We reproduce the one-loop network in figure 12b with a relabeling of the tensors that is slightly more convenient for our current purposes. This process of iteration is functionally almost identical to the inductive procedure for constructing tree networks detailed in section 3.2.

The bulk state assigned to the “bottom half” of the one-loop network in figure 12b, as defined in section 3, is the state comprised of the bulk tensor W along with the edge states ϕ and σ that correspond to neighboring extremal surfaces. As in 3, this state can be approximated by a tree tensor network in which the tensor W is replaced by an expression of the form

$$W^{A_3 A_4}_{A'_1 A'_2} \approx V_3^{A_3}_{A'_2 \gamma f} V_4^{A_4}_{A'_1 \bar{\gamma} \bar{f}} \phi^{\gamma \bar{\gamma}} \sigma^{f \bar{f}}, \quad (6.1)$$

where ϕ is a maximally entangled state on a space of dimension $e^{S_{\max}^\varepsilon(A'_1 A_4)}$, and $S_{\max}^\varepsilon(A'_1 A_4)$ is the smooth max-entropy of the bottom-half bulk state in the subregion $A'_1 A_4$. The operators V_3 and V_4 are not themselves isometries. However, when combined with the edge state ϕ and σ operators on the horizontal edges of the network, they become exact isometries flowing outwards from the newly created vertical edge.

Such a network can always be constructed for any state by entanglement distillation; however, the resulting network is only geometrically appropriate for the discretization given in figure 12a if the smooth min- and max-entropies satisfy

$$S_{\max}^\varepsilon(A'_1 A_4) = \frac{\text{Area}(\Sigma_{A_3:A_4})}{4G_N} + O\left(\frac{1}{\sqrt{G_N}}\right)$$

and

$$S_{\min}^\varepsilon(A'_1 A_4) = \frac{\text{Area}(\Sigma_{A_3:A_4})}{4G_N} + O\left(\frac{1}{\sqrt{G_N}}\right). \quad (6.2)$$

Arguments given in section 2.3 would imply equations (6.2) if the tensor network state for the bulk subregion represented by W has entropies given by the areas of extremal

surfaces in the entanglement wedge of $A_3 \cup A_4$. However, unlike the minimally entangled purification, this state has not previously been conjectured to have this property. Nevertheless, such a conjecture is closely analogous to the conjectures used in section 5, and follows intuitively from the surface-state conjecture [50]. The state defined by W and its neighboring edge states is naturally associated to the “bottom half” bulk subregion, and so the area of the extremal surface dividing A_3 and A_4 is unquestionably the “natural” geometric quantity that would be associated to the smooth min- and max-entropies of this state.

If this is indeed the case, and the bulk tensor W can be distilled into a tree network for the bottom half of the four-tensor discretization from figure 12a, then the resulting expression is a tensor network for the four-tensor discretization whose bond dimensions match the areas of extremal surfaces in the bulk. This completed network is sketched in figure 12c. The only remaining question to ask is whether the boundary state of this network accurately reproduces the original “target” CFT state.

Arguments given in section 5 imply that the one-loop network of figure 12b well approximates the original CFT state. Since the tree network for the bottom half state well approximates the original bottom half state in the one-loop network of figure 12b, the final network should still well approximate the boundary CFT state, so long as the error induced in the bottom half state from its approximation as a tree network doesn’t dramatically increase in size when the rest of the network is added.

Of course, here we run into something of a problem. The extension map upwards from the bottom half state is not an exact isometry. Indeed, as we discussed in section 5.2, it is not even an approximate isometry in the sense of section 3.2; instead we proposed that it should be called a “moral” isometry. We therefore will not have good control over the total accumulated error between the state produced by the four-tensor network and the target state. This is in contrast to the networks in section 3 (and section 5), where we could precisely control the total error that could be accumulated, so long as the RT surfaces were added in an appropriate order.

On the other hand, the edge states on the horizontal RT surface and the top half state are both normalized quantum states. This means that the upwards flowing map preserves the norm of the edge states on the horizontal RT surface, which are fully entangled. Hence we can be relatively hopeful that a generic perturbation to the bottom half state will not be dramatically blown up in size by the upwards flowing map, and the final state produced by the network in figure 12c should approximately reproduce the original CFT state.

Another way of seeing that the four-tensor network should approximately reproduce the correct boundary state is to compare the exact isometry W in figure 12b to the final downwards-flowing extension map made up of tensors V_3 and V_4 and bottom-half edge states $|\phi\rangle$ and $|\sigma\rangle$ in figure 12c. For the sake of this argument, let the full downwards-flowing extension map be denoted by B . Since the network in figure 12c was obtained from the network in figure 12b by tree network distillation, B and W must have approximately the same action on the reduced density matrices of the $|\phi\rangle$ and $|\sigma\rangle$ states on the horizontal RT surface, i.e.,

$$\|(B - W)\phi_1\sigma_1\phi_2\sigma_2\|_2 \ll 1. \tag{6.3}$$

The condition for the four-tensor network to approximately reproduce the original holographic state is for B and W to approximately agree on the *full* state for the top half of the four-tensor network, i.e.,

$$\|(B - W)\rho_T^{1/2}\|_2 \ll 1, \quad (6.4)$$

where ρ_T is the truncated tensor network state for the top half of figures 12b and 12c. These conditions are inequivalent, as we do not expect $\rho_T^{1/2}$ to be a product state across the two halves of the horizontal RT surface, because the upwards-flowing extension map is only a moral isometry. However, our version of the holographic entanglement of purification conjecture implies that the correlations between these two halves are subleading in G_N in the sense of the mutual information. Equation (6.3) can be interpreted as taking a weighted average of $(B - W)$ over the product of the marginal distributions on each half of ρ_T , while equation (6.4) can be interpreted as taking a weighted average over the joint distribution. We expect that, barring unlikely disasters, equation (6.3), which follows from conjectures given in section 5, should imply equation (6.4).

We will revisit the question of exact, approximate, and moral isometries in section 7.3, where it is shown that any geometrically appropriate tensor network of the form shown in figure 12c *must* have some moral isometries, as constructing such a tensor network with stronger bulk-to-boundary isometry conditions is inconsistent with the dynamics of the original CFT state on the boundary.

6.2 Arbitrarily fine discretizations

It is fairly easy to extend this construction to tensor networks on arbitrarily fine grid discretizations of the bulk. To iterate the one-loop (now n -loop) network, a grid is chosen like the one in figure 13, where the horizontal surfaces are extremal surfaces, and each vertical segment is an extremal surface linking the horizontal RT surfaces on either of its endpoints. The vertical surfaces are chosen via a “top-to-bottom” inductive procedure, where the top point of each segment is fixed at the bottom of the previous segment, while the bottom point is chosen to minimize the total area of the surface. In other words, each vertical segment is the minimal surface connecting neighboring horizontal RT surfaces subject to the constraint that it must continue the vertical segment above it.

To extend the one-loop iteration procedure to this discretization, we begin by constructing an n -to-one network for the top “row” of the grid, the bottom tensor of which is then distilled into a grid network for the remainder of the bulk by induction. By assuming that the bulk state represented by the bottom tensor satisfies the surface-state correspondence, and hence the holographic entanglement of purification conjecture, this top-row distillation procedure can be repeated until a network is produced for the entire grid. Note that because the bulk spacetime is curved, the grid on which we discretize cannot have right angles everywhere: in a negatively-curved spacetime, there cannot exist a quadrilateral, bounded by geodesics, with each corner having an angle of $\pi/2$. The iterative procedure for constructing a tensor network on this grid will generally produce vertical surfaces that have kinks as they pass through each horizontal RT surface.

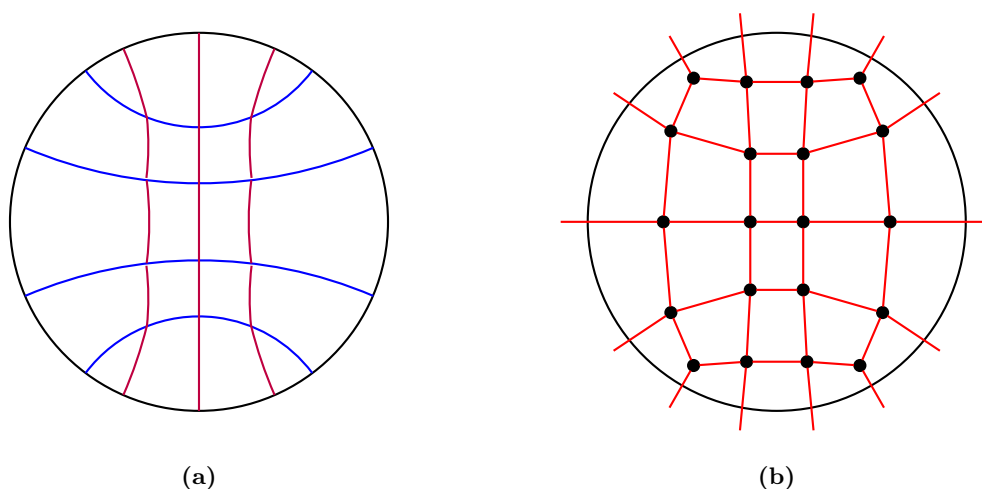


Figure 13. (a) A grid discretization for vacuum AdS_3 , for which one can find a tensor network by iterating an “ n -to-one” loop network from the top to the bottom of the grid. As explained in the text, not all surfaces in this discretization will meet at right angles. If the network is constructed “top-to-bottom”, then the bottom of each vertical segment will meet the corresponding horizontal RT surface at a right angle. (b) The dual graph of this discretization, which forms the underlying geometry for a holographic tensor network.

It is worth noting that this procedure can be used to generalize the n -to-one networks of section 5 to higher-dimensional spacetimes. In section 5, the holographic entanglement of purification conjecture was used to construct a tree tensor network for the entanglement wedge of a boundary region that was partitioned as $A_1 : \dots : A_n$. This tree tensor network was then extended to the global spacetime by an isometry on the purifying space. Generalizing this procedure to higher dimensions requires constructing a tensor network for a boundary partition that is a *grid*, not simply a one-dimensional chain. Such a partition cannot be represented by a tree tensor network, and will generally require a grid network that looks more like the one sketched in figure 13. Since we now know how to construct grid networks in $2+1$ dimensions by iteration, however, we are able to construct n -to-one networks in higher-dimensional spacetimes; one simply constructs the minimally entangled purification for a “chain” partition $A_1 : \dots : A_n$, then distills the other grid directions using the iterative techniques explained above. By repeating this procedure inductively, it is possible to construct a tensor network for a grid discretization in spacetimes of arbitrary dimension.

All our constructions in this section, as well as those in section 5, are built on the holographic entanglement of purification conjecture, which itself has yet to be proven and may not be exactly true as stated. Moreover, even if the holographic entanglement of purification conjecture is itself valid, the various generalizations of it that we used in this section could be one step too far. However, we believe that the second possibility is considerably less likely than the first. All of our constructions follow from the same basic guiding principle as the holographic entanglement of purification conjecture itself: that there should exist a state associated to any convex bulk surface (the “surface-state correspondence” [50]) and that by minimizing an entropy or a sum of entropies over all

possible purifications of a reduced density matrix, we can obtain something close to the state associated to the surface that minimizes the corresponding area or sum of areas.

Having completed our description of the details of our general construction, we offer two final motivations for believing that it is the most natural way to construct a geometrically appropriate tensor network, assuming one exists. The first is that radial flow in AdS/CFT has long been understood to be a form of renormalization group flow for the boundary CFT state. Since renormalization of a state is best understood as a process of disentangling and removing redundant degrees of freedom [7, 8, 19], the radial flow of a boundary subregion to a Ryu-Takayanagi surface should correspond to performing some position-dependent RG flow on the boundary. Our procedure for constructing holographic tensor networks consists of disentangling and discarding as many degrees of freedom as possible on a boundary subregion without changing the reduced density matrix on the complementary subregion of the boundary. If renormalization group flow is a philosophically correct approach to describing the bulk in AdS/CFT, then our constructions should also be valid, with the additional benefit that they have the potential to work at sub-AdS scales.

A second supporting motivation for our construction is that when we construct tensor networks through entanglement minimization and distillation, the resulting networks have bulk legs that are “as small as possible” while still approximately preserving the boundary state. The usual Ryu-Takayanagi inequalities for tensor networks [7] imply that the bond dimensions of a tensor network have lower bounds that are determined by the bulk geometry. If *any* geometrically accurate tensor network exists for a given CFT boundary state, then it should be found by a maximally efficient minimization procedure; we believe that our minimization procedure is the most obvious one to consider.

7 Quantum geometry

In this section we consider the effects of quantum fluctuations of the spacetime geometry, which we mentioned briefly in section 3.1 when discussing subleading entanglement in our tensor networks. We first argue that these fluctuations are best understood as quantum superpositions of tensor networks. We then point out the existence of a quantum uncertainty relation between the areas of intersecting holographic entropy surfaces, whereby a very precise measurement of the area of one such surface causes the area of the other surface to grow dramatically in size. This poses some issues for the interpretation of tensor networks whose underlying bulk discretizations contain intersecting Ryu-Takayanagi surfaces, such as those constructed in section 6. We close the section by formalizing these issues in the form of a no-go theorem that limits the isometry conditions that can be imposed on the bulk-to-boundary maps of such networks.

7.1 Superpositions of tensor networks: sweeping away the “cobwebs”

In this paper, we have generally aimed to describe a holographic boundary state with a *single* tensor network that was expected to capture the bulk geometry. In a full theory of quantum gravity, however, the bulk geometry itself is expected to be quantum mechanical, and therefore subject to quantum fluctuations around some semiclassical background. In

section 3.1, we proposed that the necessity of including in our networks some subleading edge states $|\sigma\rangle$, which are not maximally entangled, is intimately related to the existence of these fluctuations. Specifically, we argued that the $O(1/\sqrt{G_N})$ log rank of these states suggests that they correspond to fluctuations in the areas of extremal surfaces in AdS/CFT.

Another approach would be to have a single tensor network encode a single, non-fluctuating bulk geometry. In this interpretation, a holographic boundary state should be described not as a single tensor network that encodes the quantum fluctuations in the geometry, but as a weighted quantum superposition of networks, each of which describes a (very slightly different) non-fluctuating bulk geometry. The idea that fluctuations over different geometries correspond to taking a quantum superposition of different tensor networks has been previously discussed in [29, 30, 51, 52].

We can replace our single network with a superposition of tensor networks by reinterpreting its subleading tensors. In section 3.1, we suggested that because of their relatively small dimension, the σ -legs could be thought of as thin “cobwebs” adhering to the thick “girders” of the main network of ϕ -legs.³² Instead of thinking of the cobwebs as part of the tensor network, however, we can choose to interpret them as determining the weights with which the many different “girder-only” networks are superposed against one another. In doing so, one would eliminate the cobwebs associated with geometric fluctuations in any single network, instead using them to weight a superposition of “fixed” geometries. One major advantage of this approach, as we will see, is that the holographic state can be described, with high accuracy, by a superposition of only $O(1/\sqrt{G_N})$ fixed-geometry tensor networks; this is a huge improvement over the $e^{O(1/\sqrt{G_N})}$ -rank Hilbert space of fluctuations that was necessary in section 3.1.

To be more precise, consider an edge state $|\sigma\rangle$ in a holographic tensor network. By measuring this state in its Schmidt basis, we obtain a tensor network with no cobwebs on the corresponding edge. If we measure all cobwebs in the network according to this procedure, then the resulting network for any measurement outcome has no cobwebs on any edge. To write the original holographic state in terms of these measured networks, we need to use a superposition of the networks associated with every possible measurement outcome, with each network weighted by the corresponding Schmidt coefficient in each measured cobweb state $|\sigma\rangle$. Given the absence of the cobwebs that we associated with fluctuations in geometry, one might reasonably suggest that each of these tensor network corresponds to a fixed, non-fluctuating geometry. The original, holographic state can then be written as a superposition of appropriately weighted geometries.

So far, this “superpositions” framework is merely a different way of interpreting the tensor networks we already constructed in section 3. However, interpreting the full, semiclassical tensor network as an ensemble of “girder-only” networks with slightly different geometries suggests that we should allow the dimensions of the girders to vary for different networks in the superposition. This can be accomplished with a slight adjustment to the block-averaging procedure from section 3.1, in which we allow different blocks to have

³²In such an approach it would be natural to represent long range entanglement of geometry fluctuations using cobwebs that have legs extending nonlocally to many different girders.

different widths. In fact, we will now show that for any fixed error ε , the “target” holographic state can be approximated to within tolerance ε by choosing the block widths to be proportional to $\varepsilon e^E O(\sqrt{G_N})$, where E is the block-averaged eigenvalue of the modular Hamiltonian $K = -\log \rho$. This makes the total number of blocks, and hence the total number of girder-only networks in the superposition, order $O(1/\sqrt{G_N})$.

To define a general block-averaging prescription, let n be an index that labels the eigenvalues of K as in section 3.1, and let w_n be the width of the block containing the n^{th} eigenvalue. If $p_n = e^{-E}$ are the eigenvalues of the density matrix ρ , then the one-norm error induced by replacing each p_n with the average eigenvalue within its block is given by

$$\varepsilon = \int |p_n - \bar{p}_n^\varepsilon| dn \approx -\frac{1}{4} \int \frac{dp_n}{dn} w_n dn, \tag{7.1}$$

where we have approximated p_n as being roughly linear within each block and approximated the index n by a smooth function. To find the optimal block-averaging procedure for representing ρ , we want to minimize the number of blocks — and hence the total number of tensor networks that must be superposed to describe ρ — subject to the constraint of a fixed error ε . The total number of blocks is given by

$$N_{\text{blocks}} = \int \frac{1}{w_n} dn, \tag{7.2}$$

and so the function w_n that minimizes the total number of blocks subject to the constraint given by (7.1) satisfies

$$\frac{1}{w_n^2} = \frac{\lambda}{4} \frac{dp_n}{dn}, \tag{7.3}$$

where λ is a Lagrange multiplier that does not depend on n . We may solve for the value of λ by plugging this expression back into (7.1), yielding

$$\sqrt{-\lambda} = \frac{1}{2\varepsilon} \int \sqrt{-\frac{dp_n}{dn}} dn. \tag{7.4}$$

With respect to the smooth index n , the density of states is given by

$$D(E) = \frac{dn}{dE}, \tag{7.5}$$

and so the eigenvalues $p_n = -e^{-E}$ of the density matrix satisfy

$$\frac{dp_n}{dn} = -\frac{e^{-E}}{D(E)}. \tag{7.6}$$

It follows from (7.4), then, that the Lagrange multiplier λ is given by

$$\sqrt{-\lambda} = \frac{1}{2\varepsilon} \int \sqrt{D(E)e^{-E}} dE. \tag{7.7}$$

Using expression (7.3) for the optimal block widths w_n , we find that the optimal number of blocks is given by

$$N_{\text{blocks}} = \int \frac{1}{w_n} dn = \frac{1}{4\varepsilon} \left[\int \sqrt{D(E)e^{-E}} dE \right]^2. \tag{7.8}$$

In section 2.3, we argued that the spectrum of K is tightly constrained around the leading saddle point in $D(E)e^{-E}$. Near this saddle point, the function can be approximated by a normalized Gaussian of width $O(1/\sqrt{G_N})$, i.e.,

$$D(E)e^{-E} = O(\sqrt{G_N})e^{-O(G_N)(E-S)^2}. \tag{7.9}$$

The optimal number of blocks (7.8) may therefore be computed as

$$N_{\text{blocks}} = \frac{1}{\epsilon} \frac{1}{O(\sqrt{G_N})}, \tag{7.10}$$

as we claimed above. The optimal block widths, w_n , satisfy

$$w_n = \frac{2}{\sqrt{-\lambda}} \sqrt{-\frac{dn}{dp_n}} = \epsilon O(G_N^{1/4}) \sqrt{D(E)e^E}. \tag{7.11}$$

Near the saddle point, the variation in the right hand side of (7.9) is subleading and so we have

$$D(E) \approx O(\sqrt{G_N})e^E. \tag{7.12}$$

It follows that the block widths in the optimal block-averaging procedure are proportional to e^E . Assuming that the modular energy E has a holographic interpretation as $A/4G_N$, where A is the area of the minimal surface, this is the right size for the tensor network geometry to match the semiclassical geometry of the holographic state.

Unlike our construction in section 3, however, the optimal block-averaging procedure represents the semiclassical holographic geometry as a superposition of $O(1/\sqrt{G_N})$ rather than $e^{O(1/\sqrt{G_N})}$ “fixed-geometry” networks. We take this as a suggestion that the superposition framework, in which one allows the different tensor networks in the ensemble to have slightly different “fixed” geometries, is a more efficient and informative description of the holographic state.

In this superposition framework, it is natural to interpret a single tensor network in the superposition as corresponding to an (approximate) eigenstate of the bulk area operator [29, 30, 53, 54]. If the usual bulk state generated by a path integral corresponds to a canonical ensemble of the area operator, then the state of a single tensor network corresponds to a microcanonical ensemble which takes values over a tiny range of areas.³³ This interpretation is compelling because it makes the flat entanglement spectrum of many tensor network models (in which all Rényi entropies are equal) into a feature rather than a bug. In a single tensor network, the Rényi entropy *is* flat; to get a state with a non-flat spectrum, one must take superpositions of different geometries.

7.2 Uncertainty relations for intersecting Ryu-Takayanagi surfaces

Having said all this, we will now identify a serious issue with the approach of [29, 30]. Namely, area-eigenvalue states with a flat or nearly-flat entanglement spectrum do not

³³We do not expect that the area spectrum will contain large exact degeneracies, but because of the ϵ -smoothing, we can approximate an area spectrum with exponentially small gaps with a degenerate spectrum.

correspond to a bulk geometry that is similar to the original state. As a result, it is not possible to construct geometrically appropriate tensor networks that are in a microcanonical area ensemble for multiple directions simultaneously.

The problem is that in general relativity, the area and boost angle are canonically conjugate quantities obeying a Heisenberg uncertainty relation [55, 56]:

$$\Delta E \Delta t \geq \frac{1}{2}, \tag{7.13}$$

where the modular Hamiltonian is $E = \text{area}/4G_N + O(1)$ [57] and the conjugate time t is the boost angle (in hyperbolic radians). Obtaining a flat entanglement spectrum in a holographic state requires constraining ΔE to a small, $O(1)$ number, and hence measuring the area of the corresponding surface to within a tolerance of $O(G_N)$.

Since a single bin in our construction measures E to an accuracy of $O(\epsilon)$, (7.13) implies that the uncertainty in the corresponding boost angle is $O(1/\epsilon)$. This is quite large, and in fact it is large enough to take us out of the validity of the static slice regime. In particular, introducing a large crease of extrinsic curvature at the horizontal Ryu-Takayanagi surface of figure 12c will make it so that the area of the vertical HRT surface (which follows a *spacetime* geodesic and therefore no longer lies on the creased slice) will have a significantly greater area than the minimal surface on the original static slice, as shown in figure 14.³⁴ Since this is true for either sign of the boost angle, the expectation value of the area given an uncertain boost is also larger.

This uncertainty principle implies that if we start by distilling information on the horizontal surface, consider *just one term* of the resulting superposition, and then attempt a “vertical” distillation, then the leading $O(1/G_N)$ part of the vertical entropies will be larger than on the original static slice. In other words, it will not be possible to construct a geometrically appropriate tensor network for a single term of the superposition.³⁵ This problem is related to the issues for dynamical tensor networks that we discuss in section 8.4.

7.3 A no-go theorem for the four-tensor network

The uncertainty relationship discussed in the previous section shows that we cannot simultaneously measure the areas of horizontal and vertical Ryu-Takayanagi surfaces with high accuracy. This might make one wonder whether it is really possible, in a single tensor network, to localize information on both the vertical and horizontal Ryu-Takayanagi surfaces simultaneously, as required for our iterative constructions in 6 to be geometrically accurate.

In this section, we will prove that there is indeed such an obstruction preventing the construction of certain kinds of geometrically accurate tensor networks with crossing Ryu-Takayanagi surfaces, even in the simplest case of a four-tensor network (section 6.1). However, our no-go theorem is only valid for tensor networks having approximate isometry

³⁴In evaluating the entropy of the HRT surface, it is helpful to use a boost-invariant UV cutoff surface, so that the area of surfaces on the slice is independent of the boost angle.

³⁵In the exact AdS/CFT bulk state, this spuriously large entropy must disappear when we take all terms in the superposition, due to destructive interference. However, it can be difficult to keep approximations under control when there is destructive interference among a large number of terms, since small errors can accumulate and leave a substantial remainder.

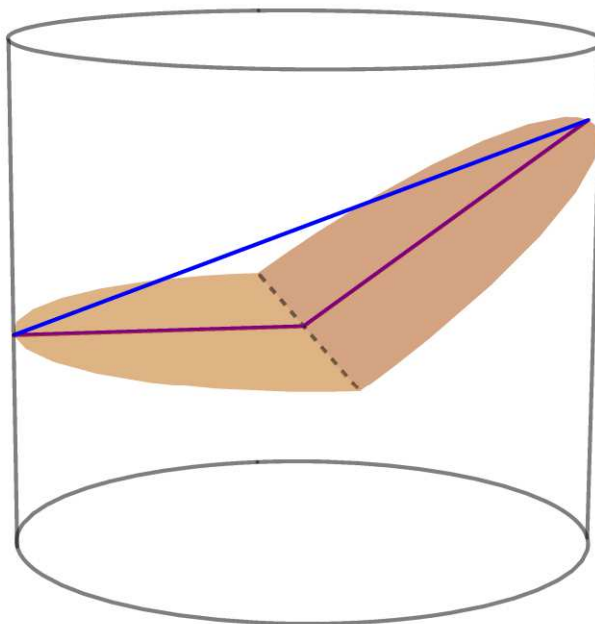


Figure 14. A spacelike slice of vacuum AdS_3 formed by boosting one half of the $t = 0$ static slice. This boost introduces a nontrivial extrinsic curvature at the dashed line in the surface. The original RT surface for a particular boundary region on the $t = 0$ slice, sketched by a purple line, is no longer the entangling surface for the corresponding boundary region in the half-boosted state. Instead, one must consider the HRT surface, sketched in blue, which has strictly larger area.

properties that are stronger than those of our actual network. Thus, we can optimistically hope that our tensor network constructions are still valid.³⁶

When constructing sub-AdS scale tensor networks in sections 5 and 6, we found that the bulk-to-boundary “extension maps” associated with a particular network were not generally exact isometries. In the n -to-one loop networks of section 5, at least the “downward-flowing” map could be shown to be an exact isometry — or, depending on the order in which the “top-half” RT surfaces were distilled in constructing a tree network for the MEP, at least an approximate isometry in the sense of equation (3.25). In the full sub-AdS network of section 6, however, we found that the bulk-to-boundary maps in our network could *not* generally be proven to be exact or approximate isometries, and were in fact only “morally” isometric in the sense that they preserved the normalization of the state and preserved the entanglement entropy to leading order. The non-exactness of these isometries stands in contrast to the holographic tensor network toy models of AdS/CFT introduced in [9].

In this section, we prove that geometrically appropriate tensor networks for generic bulk discretizations of static states in AdS/CFT *cannot*, in fact, have bulk-to-boundary extension maps which are all exact or even approximate isometries. We formalize this for the four-tensor network for vacuum AdS_3 of section 6 in the following theorem:

³⁶If not, then we believe the construction in 5, where we localize information on a single Ryu-Takayanagi surface, will still be valid.

Theorem 1. *No four-tensor network of the form shown in figure 12c can simultaneously satisfy the following four properties:*

- (i) *Each leading edge state $|\phi\rangle$ is maximally entangled on a Hilbert space of dimension $e^{S\pm o(S)}$, where S is proportional to the area of the corresponding bulk surface, satisfying*

$$S = \frac{\text{area}}{4G_N} + o\left(\frac{1}{G_N}\right). \quad (7.14)$$

- (ii) *Each subleading edge state $|\sigma\rangle$ is submaximally entangled on a Hilbert space of dimension $e^{o(S)}$.*

- (iii) *The four extension maps that map either side of either RT surface to the boundary are all approximate isometries in the sense that*

$$V^\dagger V \approx \mathbb{1} \quad (7.15)$$

for any such map V , where this approximation means that $V^\dagger V$ is close to the identity in the operator norm.

- (iv) *The boundary state of the tensor network approximately reproduces the boundary state of the AdS_3 vacuum with high fidelity.*

Suppose that the four-tensor network shown in figure 12c *does* satisfy all conditions given in Theorem 1. We denote the “upwards-pointing” extension map from the horizontal RT surface as T , and the “downwards-pointing” map as B . In other words, if we were to collapse all the tensors in the top and bottom halves of figure 12c, excluding the tensors on the horizontal RT surface itself, then the resulting tensor network would have only T as its top-half tensor and B as its bottom-half tensor (sketched in figure 15a). We denote the actual CFT state as $|\psi^{\text{CFT}}\rangle$, and the tensor network state as $|\psi^{\text{TN}}\rangle$. Assumption (iv) of Theorem 1 ensures

$$|\psi^{\text{CFT}}\rangle \approx |\psi^{\text{TN}}\rangle. \quad (7.16)$$

From the form of the tensor network shown in figure 15a, we see that $|\psi^{\text{TN}}\rangle$ has the form

$$|\psi^{\text{TN}}\rangle = (T \otimes B)|\phi\sigma\rangle, \quad (7.17)$$

where $|\phi\sigma\rangle$ represents the combined pure state of all edge states $|\phi\rangle$ and $|\sigma\rangle$ on the horizontal RT surface.

Let K_B^{CFT} be the modular Hamiltonian of the bottom-half boundary state $\psi_{A_3A_4}^{\text{CFT}}$, and $e^{iK_B^{\text{CFT}}}$ the corresponding boost operator. Since this operator is unitary, we have

$$e^{iK_B^{\text{CFT}}} |\psi^{\text{CFT}}\rangle \approx e^{iK_B^{\text{CFT}}} |\psi^{\text{TN}}\rangle. \quad (7.18)$$

Since the reduced state of the tensor network on A_3A_4 approximately reproduces the reduced state of the CFT on the same region, their modular Hamiltonians approximately

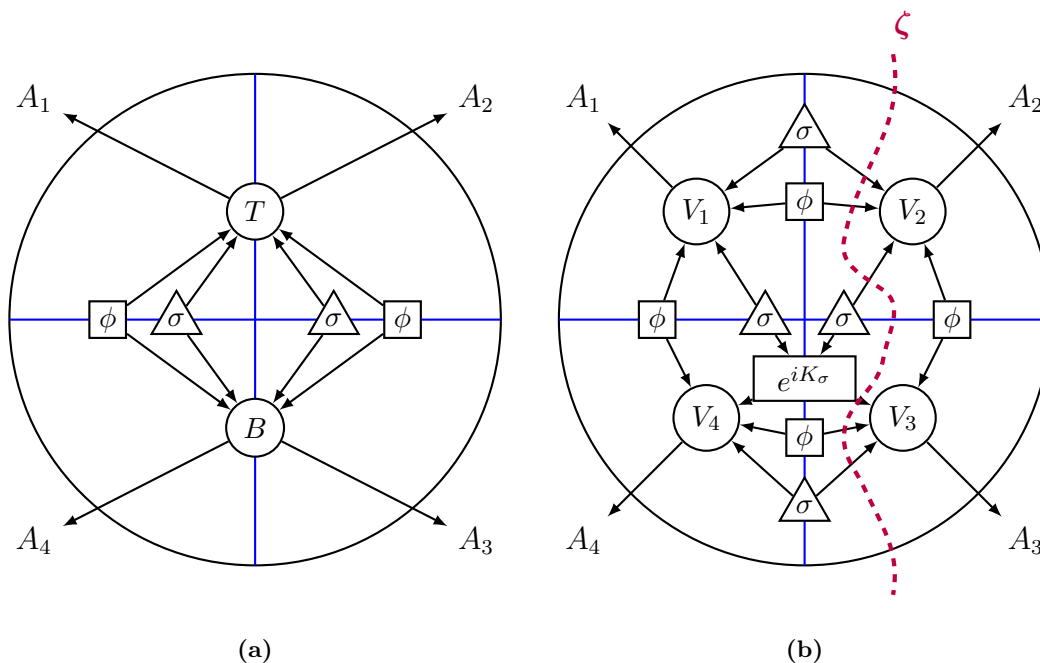


Figure 15. Figures to accompany the no-go theorem for the four-tensor network. (a) The extension maps for the top and bottom halves of the horizontal RT surfaces are drawn explicitly as tensors T and B . (b) The boost operator on the “bottom half” of the boundary, A_3A_4 , can be represented on the tensor network by a unitary operator e^{iK_σ} that acts only on the subleading edge states $|\sigma\rangle$. The graph cut ζ , sketched here, has dimension given to leading order by the combined size of the $|\phi\rangle$ Hilbert spaces on the vertical RT surface.

agree.³⁷ It follows that the action of the modular Hamiltonian on the CFT state can be represented in the tensor network as

$$e^{iK_B^{\text{CFT}}} |\psi^{\text{CFT}}\rangle \approx e^{iK_B^{\text{TN}}} |\psi^{\text{TN}}\rangle, \quad (7.19)$$

where K_B^{TN} is the modular Hamiltonian of the bottom-half boundary state $\psi_{A_3A_4}^{\text{CFT}}$ in the tensor network.

We see immediately from the form of the tensor network state given in equation (7.17) that this modular Hamiltonian takes the explicit form

$$K_B^{\text{TN}} = -\log \psi_{A_3A_4}^{\text{TN}} = -\log \text{Tr}_{A_1A_2} (TB\phi\sigma B^\dagger T^\dagger). \quad (7.20)$$

Condition (iii) of Theorem 1 ensures that $T^\dagger T$ is close to the identity, ensuring that the partial trace over A_1A_2 in the above expression can be replaced by a partial trace over the domain of T , which we call $\mathcal{H}_f \otimes \mathcal{H}_\gamma$. Here, as in section 3, \mathcal{H}_f corresponds to the subleading states $|\sigma\rangle$ while \mathcal{H}_γ corresponds to the maximally entangled states $|\phi\rangle$. The modular Hamiltonian of the reduced tensor network state on A_3A_4 therefore satisfies

$$K_B^{\text{TN}} \approx -B(\log \text{Tr}_{f\gamma} \phi\sigma) B^\dagger. \quad (7.21)$$

³⁷The modular Hamiltonians may not actually be close in the sense of the operator norm; however, their action on the global states $|\psi^{\text{TN}}\rangle$ and $|\psi^{\text{CFT}}\rangle$ are approximately the same.

In other words, the modular Hamiltonian of the tensor network on the bottom-half boundary region A_3A_4 can be approximately represented by

$$K_B^{\text{TN}} \approx BK_{\phi\sigma}B^\dagger, \tag{7.22}$$

where $K_{\phi\sigma}$ is the modular Hamiltonian of the reduced RT surface-state $\text{Tr}_{f_\gamma} \phi\sigma$.

Since the edge states $|\phi\rangle$ are maximally entangled, their modular Hamiltonian when restricted to either side of the RT surface is simply a multiple of the identity. This contributes only a normalization factor to the overall network. We therefore write the modular Hamiltonian of the tensor network on A_3A_4 as

$$K_B^{\text{TN}} \approx BK_\sigma B^\dagger, \tag{7.23}$$

up to an additive constant coming from the normalization, which we ignore. Here K_σ is an operator that acts only on the subleading boundary states $|\sigma\rangle$, and only acts on one side of the RT surface (in this case, the bottom half).

Returning to the tensor network expression for the modular flow of the CFT given in equation (7.19), we see that the boost operator on the bottom half of the boundary CFT state can be represented on the tensor network as

$$e^{iK_B^{\text{CFT}}} |\psi^{\text{CFT}}\rangle \approx Be^{iK_\sigma} B^\dagger |\psi^{\text{TN}}\rangle. \tag{7.24}$$

From the expression for the tensor network state given in (7.17), we may rewrite this expression as

$$e^{iK_B^{\text{CFT}}} |\psi^{\text{CFT}}\rangle \approx (T \otimes B)e^{iK_\sigma} |\phi\sigma\rangle, \tag{7.25}$$

where we have used condition (iii) of Theorem 1 to ensure that $B^\dagger B$ is close to the identity. This final tensor network representation for the boosted CFT state is sketched in figure 15b.

Equation (7.25) essentially tells us that the modular flow of the CFT on the boundary region A_3A_4 can be represented by a unitary operator that acts only on the subleading states $|\sigma\rangle$ on the horizontal RT surface. This conclusion, however, contradicts assumptions (i) and (ii) of Theorem 1, which restrict the bond dimensions of the network. To see this contradiction, consider the entanglement entropy of the boundary region A_2A_3 in the CFT state $e^{iK_B^{\text{CFT}}} |\psi^{\text{CFT}}\rangle$. In the bulk, this operator acts as a boost on the entanglement wedge of A_2A_3 . In vacuum AdS_3 , it is easy to show that the entangling surface of A_2A_3 in the boosted state has greater area in the boosted state than in the unboosted state. By the HRT formula [4], it follows that the entanglement entropy of A_2A_3 in the boosted state must be greater than the corresponding entropy in the unboosted state at leading order, i.e., $S(A_2A_3) > \text{area}(\text{vert})/4G_N$, where $\text{area}(\text{vert})$ represents the area of the vertical RT surface (see figure 14).

However, this is in direct contradiction a bound derived by Swingle in [7]! There it was shown that the entanglement entropy of a boundary subregion A in any tensor network is bounded above by $\log \dim \zeta$ for any graph cut ζ that partitions A from its complement. From the tensor network representation of the boosted state given in equation (7.25) and

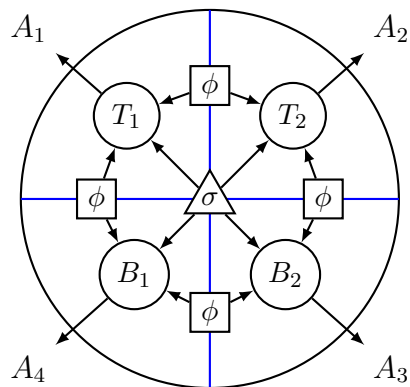


Figure 16. A four-tensor network with explicit entanglement between the right and left sides of the horizontal RT surface, where all four tensors share a common cobweb state. Adding subleading entanglement like this cannot be used to subvert our no-go theorem.

the graph cut sketched in figure 15b, conditions (i) and (ii) of Theorem 1 imply that the entanglement entropy of A_2A_3 in the boosted state satisfies

$$S(A_2A_3) \leq \frac{\text{area}(\text{vert})}{4G_N} \tag{7.26}$$

to leading order in G_N . The bound given in (7.26) clearly contradicts the HRT formula for the boosted state.

We conclude that no geometrically appropriate four-tensor network can be constructed for the AdS_3 vacuum with approximate isometries from each RT surface to the boundary. This provides partial justification for our relatively weak isometry conditions: both the state-dependent approximate isometry condition of section 3 and, in particular, the “moral” isometries of sections 5 and 6. While one might initially think that stronger isometry conditions should be possible in a tensor network construction of a holographic state, it turns out that such conditions are incompatible with the dynamics of AdS/CFT. The geometrically appropriate four-tensor network constructed in section 6 avoids our no-go theorem precisely *because* some of its extension maps are only moral isometries.³⁸

Note that while our proof of the no-go theorem required the existence of a subleading state on the horizontal RT surface to absorb the action of the modular Hamiltonian, it did not necessarily require that the subleading states take the exact form shown in figure 12c. In fact, the same proof would apply for many different kinds of cobweb entanglement (using the terminology of section 7.1, where we referred to maximally entangled states as “girders” and subleading states as “cobwebs”), so long as the cobweb states are able to absorb the modular Hamiltonian on the horizontal RT surface and do not interfere with the size of the girders. In particular, one could consider a four-tensor network with a generic cobweb state that is entangled among all four quadrants of the network, sketched in figure 16. Such a tensor network still satisfies all the necessary conditions of our no-go theorem, and

³⁸A more pessimistic interpretation would be that either the entropy of purification conjecture, or one of the additional conjectures that we made in section 6, fails to hold in the form necessary for the constructions in section 6 to go through.

so cannot have approximate bulk-to-boundary isometries while accurately reproducing the dynamics of the boundary CFT state.

8 Discussion

In this article, we have given a constructive procedure for distilling the spatial geometry of a static spacetime from the quantum entanglement of a boundary CFT, including structure at sub-AdS scales. We described a general procedure for constructing large-scale “tree tensor networks” (section 3), explained how the holographic entanglement of purification conjecture could be used to partition the information on a single Ryu-Takayanagi surface below the AdS scale (section 5), and proposed an iterative approach for constructing a sub-AdS tensor network by distilling the geometry incrementally (section 6). We showed that tree tensor networks are always quantum error correcting in the appropriate sense (section 4), and suggested that our more general, sub-AdS constructions should have similar properties. We also proved an important no-go theorem (section 7) showing that the bulk-to-boundary maps of any geometrically appropriate tensor network for AdS/CFT cannot be exact or approximate isometries in the usual sense.

While our constructive procedures can always be performed so long as the smooth min- and max-entropies of various surface subregions agree at leading order, our holographic conjectures ensure that the resulting tensor network geometry also matches the geometry of the corresponding AdS/CFT state in the sense that the Hilbert space dimensions of its legs match the areas of corresponding spacetime surfaces. This constitutes significant data about the spacetime geometry, in the sense that a sufficiently large set of bulk area observables can be used to reconstruct the metric, as has been shown explicitly in four spacetime dimensions [58] and likely holds more generally. Furthermore, our holographic conjectures ensure that any two choices of iterative construction for a single bulk discretization should produce tensor networks that differ only at subleading orders in $1/G_N$; in the limit as the discretization scale is taken arbitrarily small, any two tensor networks corresponding to *different* discretization schemes should converge up to subleading corrections. So if our conjectures hold, then for the first time we have used tensor networks to obtain the “it” of continuous spatial geometry from the “qubit” of quantum entanglement for states in full AdS/CFT.

In fact, our dictionary for constructing a tensor network uses *only* quantum information properties of the boundary state. This implies that, if we regulate the CFT on a lattice, acting on any lattice points with local unitaries does not change the resulting bulk geometry. This raises some philosophical puzzles given that entanglement is not a linear quantum observable and therefore cannot be measured by a normal quantum observation [59]. These puzzles may be related to the AMPS firewalls paradox [60–63] and claims that the construction of geometry is necessarily state-dependent [64, 65], but in this article we make no claims about which geometrical features are truly measurable.

8.1 Entanglement shadows

In describing our constructions, we usually had in mind either the vacuum state or small perturbations around it. As a result, we have mostly ignored some subtleties than can

appear in more complicated states. However, we anticipate that our results can be extended to essentially arbitrary spacetimes. In particular, it should be possible to use our results to probe the geometry of the so-called “entanglement shadow,” i.e., a bulk region that cannot be probed by normal Ryu-Takayanagi surfaces [66–69]. Such entanglement shadow regions appear, for example, in the region surrounding a massive star or black hole. We can probe such regions using our iterative construction in section 6, beginning by dividing RT surfaces that pass just outside the entanglement shadow into pieces using the holographic entanglement of purification. By iteratively constructing new surfaces anchored to these RT surfaces, it is clear that we can get farther into the bulk than we could get by starting on the boundary.³⁹

This raises the question of how deep into a general bulk we can probe. It is particularly interesting to ask this question in the context of wormholes extended between multiple asymptotic CFT regions [49, 71–73]. Because our construction relies on taking consecutive bipartite divisions of the system, our construction always requires starting with the full entangled state of all the asymptotic CFT regions. We can then easily construct a Hilbert space representing the compact RT surface of an entire connected boundary component. It is not clear, however, that there will always be nontrivial entanglement wedge cross-sections extending to the global RT surface, because it may be the case that the minimal cross-section of the entanglement wedge will close off on the boundary rather than traveling down the wormhole throat.

In light of the above reflections, we believe that the only possible obstruction to continuing deeper into the bulk using our iterative procedure would be a locally minimal area surface Σ for an entire connected component of the boundary with the property that for any bipartite division of $\Sigma = \Sigma_1 \cup \Sigma_2$, the minimal area surface separating Σ_1 and Σ_2 is always whichever of Σ_1 or Σ_2 has the least area. Such a surface Σ would be analogous to a Haar-random pure state, for which the entanglement entropy of any bipartite division scales as the volume of the smaller subsystem. Even in these cases, however, one might sometimes be able to go deeper if there exists a nontrivial entanglement wedge cross-section from Σ to the true global minimum. It would be interesting to confirm this intuition using geometrical proofs similar to [68].

8.2 Bit threads

Another potential avenue for future work is the apparent connection between our construction and the bit-thread formalism for holographic entanglement [24]. Bit threads describe boundary entanglement in terms of smooth flows on the bulk geometry, with intuition inherited from the machinery of maximal flows on discrete graphs. In some sense, then, the bit thread formalism is inspired by the notion of discretizing the bulk geometry of a holographic state; discretizing the bulk geometry of a holographic state while preserving its entanglement structure is exactly what we have done in this paper by constructing tensor networks for AdS/CFT. Since our construction is largely based on the Ryu-Takayanagi for-

³⁹For a similar discussion of how entanglement shadows can be probed by extremal surfaces anchored to points in the bulk, see [70].

mula and its generalizations, and bit threads interpretations of holographic entanglement are formally equivalent to RT, our work could be reframed in the language of bit threads.

The bit threads formalism suffers from the fact that known holographic entropy inequalities [74, 75] seem generally more difficult to prove in the language of bit threads than in the usual Ryu-Takayanagi picture [76, 77]. It is possible, however, that the bit thread formalism is more useful than the Ryu-Takayanagi formula for describing localization of boundary entanglement to small “cells” of the bulk. Many times in this work, especially in section 3.2, we have appealed to notions of “information flow” toward or away from bulk subregions to describe our construction intuitively. This notion is best expressed in the language of bit threads, and one might expect that reframing our work in the bit thread formalism could yield new insight into the information content of bulk subregions in tensor networks with sub-AdS locality. In particular, one might ask the following: what happens when one tries to define a spacetime flow that maximizes the flow through a single bulk subregion in our tensor network construction, or a combination of bulk subregions? Does this maximal bulk flow have a natural interpretation in terms of boundary information? Flows of this type were recently discussed in [78]. We leave further analysis of this question for future work.

8.3 Fiber directions

An important aspect of the holographic bulk that cannot be decomposed in our construction is the geometry of the Kaluza-Klein fiber directions. For example, the ABJM model vacuum is dual to $AdS_4 \times S_7$ [79], while the $\mathcal{N} = 4$ Super Yang-Mills vacuum is dual to $AdS_5 \times S_5$ [1]. In such states, we cannot possibly use entropy of purification to subdivide the S_n factor into smaller pieces, because spherical symmetry implies that all entanglement cross-sections will be symmetrical.⁴⁰ In our construction, these fiber directions simply go along for the ride, without being subdivided, even though their radius of curvature is of the same order as that of the AdS factor, which we do subdivide.⁴¹

Note that, as in all theories with large extra dimensions, the 10 or 11 dimensional Planck length will be parametrically larger than the effective Planck scale of the Kaluza-Klein reduced theory on AdS. This raises the question of whether ignoring the fiber directions might actually allow us to subdivide the AdS factor at a finer scale than we could if we were also subdividing the sphere. After all, our construction depends on the Hilbert space dimension being large, and including all the modes of the sphere makes the Hilbert space larger than it would be otherwise. It seems likely that higher-curvature/stringy corrections to the Ryu-Takayanagi formula would prevent this from working, but this subject bears further investigation.

⁴⁰Interestingly, it might be possible to do better in excited geometries which are not spherically symmetric.

⁴¹See [80, 81] for a possible idea for how to understand the fiber directions in terms of entanglement between different field degrees of freedom. If this idea is correct it would be interesting to combine it with our construction to obtain a discretization of the full space.

8.4 Dynamics

Another important unanswered question is how to extend our construction to dynamical settings. The main inhibition to extending the construction to holographic states with dynamical bulk spacetimes is that while boundary entropies are still given by the areas of extremal surfaces in the bulk, those extremal surfaces can generally not all be chosen to lie in a single spacelike bulk slice [4]. In a static spacetime with Killing time parameter t , by contrast, all Ryu-Takayanagi surfaces of the $t = 0$ boundary slice can be made to lie in the $t = 0$ slice of the bulk. Such a proposal would presumably require either (i) discretizing time to produce a d -dimensional tensor network for a d -dimensional spacetime [52, 82], or else (ii) constructing different tensor networks for different Cauchy slices of the bulk [83], which nevertheless give rise to the same boundary state. In the latter case, it is tempting to identify the gauge equivalence of tensor networks with the Hamiltonian constraint of the continuum bulk general relativity, since both involve an equivalence of states at different times.

Unfortunately, it is impossible for tensor networks on dynamical Cauchy slices to be geometrically appropriate in the same sense as on static slices. Let us consider any boundary region R , and attempt to construct a tensor network on a Cauchy slice Σ which does *not* include the HRT surface X_R . By the maximin construction [84], the minimal area cut γ always has less area than the HRT surface: $\text{area}(\gamma) < \text{area}(X_R)$ (for an example, see figure 14). On the other hand, the Swingle bound [7] requires that $\text{area}(\gamma_R) \geq S(R) = \text{area}(X_R)$. This is a contradiction, unless we allow the log of the bond dimension to exceed the area even at leading order.

Perhaps, then, dynamical tensor networks are described by tensor networks that include long, nonlocal links connecting different parts of the network, which carry $O(1/G_N)$ amounts of information. Presumably we could still construct a tensor network by our minimization procedures. But for a general boundary slice, only a single family of non-intersecting HRT surfaces could be simultaneously placed on the same Cauchy slice and represented by edges with geometrically appropriate bond dimensions. The other directions would have to have information flow exceeding their area.

It may still be possible to construct such tensor networks in a compelling way by using the modular flow techniques of [85, 86]. In this picture, when two flat slices meet at an HRT surface with a nonzero boost angle, there is a nonlocal exchange of information along the RT surface due to the modular flow associated with the boost. (Otherwise, the maximin principle would be violated.) This clarifies that, although our construction localizes the information of the RT surface at sub-AdS scales, this localization must be understood as only being valid in a particular Lorentz frame of reference.

On the other hand, since the modular flow relates the dynamical and static cases, the dynamical problems posed by the Swingle bound are to some degree already present in the static case. This observation led to our no-go theorem in section 7.3, which shows that geometrically appropriate tensor networks for the AdS/CFT correspondence cannot generally be expected to have approximate bulk-to-boundary isometries when Ryu-Takayanagi surfaces intersect in the bulk.

8.5 Geometry from entanglement

The fact that our distillation procedure depends only on the entanglement structure of the state suggests that it may be more broadly applicable to other kinds of entangled states, perhaps e.g. those that are hypothesized to live on so-called holographic screens [87, 88] in cosmology. Given a quantum state on a lattice, one need only check that its smooth min- and max-entropies of purification have favorable properties; if so, it will project a holographic state onto its interior. The surface-state conjecture implies that the quadrilaterals of our tensor network grid (see e.g. figure 13a) have suitable holographic states living on their boundaries. In principle, we could use our construction to determine these boundary states explicitly.

We do not expect our quadrilaterals to be associated with “perfect tensors” [9], since cutting a quadrilateral along the diagonal results in a surface with less area than the two other sides of the triangle, allowing a nontrivial distillation to be performed along the diagonal. But we do expect that there will be an approximate isometry (in the sense of section 3.2) mapping any *one* of the edges to the other three edges. In general, it will be important to prove as many isometry-like relations as possible (subject to the constraints of our no-go theorem), both for the purposes of quantum error correction and to determine how sensitive the tensor network is to the precise order in which distillations are performed.

This article goes in the direction of starting with a boundary state and analyzing what the holographic tensor network must be. A complementary approach would be to start with the tensors associated with different kinds of geometries, and then synthesize them back together into an arbitrary geometry. In doing so it would be important to check that all expected isometries continue to hold. It would also be critical to show that, to high accuracy, the tensor network associated with a geometrical region does not significantly depend on either its external spatial context, or the methodology used to construct it.

If this can be done, then the holographic principle would finally be freed from the straitjacket of asymptotically AdS boundary conditions. It could be applied equally well to universes with other asymptotic structures, or even to closed cosmologies! In the latter case, the tensor network could be evaluated to give some complex number for each possible choice of spatial geometry. Such a “tensor network partition function” would in effect define a special cosmological state over the space of 3-metrics. It would be interesting to determine what relationship this special state might have to other proposals for special initial conditions, e.g. the Hartle-Hawking state.

Acknowledgments

We would like to acknowledge useful conversations with Ahmed Almheiri, Chris Akers, Raphael Bousso, Xi Dong, William Donnelly, Dan Harlow, Patrick Hayden, Ted Jacobson, Isaac Kim, Aitor Lewkowycz, Juan Maldacena, Don Marolf, Masamichi Miyaji, Ali Mollabashi, Xiao-Liang Qi, Dan Ranard, Eva Silverstein, Steve Shenker, Brian Swingle, Tadashi Takayanagi, and Guifre Vidal. We especially thank Patrick Hayden, Brian Swingle, and Michael Walter for sharing with us their unpublished work on smooth min- and max-entropies in quantum field theory. GP, JS, and AW were supported by the Simons Founda-

tion (“It from Qubit”), AFOSR grant number FA9550-16-1-0082, and the Stanford Institute for Theoretical Physics. AW was also supported in part by the John Templeton Foundation, Grant ID# 60933, and would like to thank the Centro Atómico Bariloche, the California Institute of Technology, and the Kavli Institute for Theoretical Physics (supported in part by NSF grant PHY-1748958) for hospitality at key stages of this project. NB is supported by the National Science Foundation under grant number 82248-13067-44-PHPXH and by the Department of Energy under grant number DE-SC0019380. He would also like to thank the SITP for hospitality while part of this work was completed. The opinions expressed are those of the authors and do not necessarily reflect the views of funding agencies.

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] J.M. Maldacena, *The large N limit of superconformal field theories and supergravity*, *Int. J. Theor. Phys.* **38** (1999) 1113 [[hep-th/9711200](#)] [[INSPIRE](#)].
- [2] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Phys. Rev. Lett.* **96** (2006) 181602 [[hep-th/0603001](#)] [[INSPIRE](#)].
- [3] S. Ryu and T. Takayanagi, *Aspects of holographic entanglement entropy*, *JHEP* **08** (2006) 045 [[hep-th/0605073](#)] [[INSPIRE](#)].
- [4] V.E. Hubeny, M. Rangamani and T. Takayanagi, *A covariant holographic entanglement entropy proposal*, *JHEP* **07** (2007) 062 [[arXiv:0705.0016](#)] [[INSPIRE](#)].
- [5] A. Lewkowycz and J. Maldacena, *Generalized gravitational entropy*, *JHEP* **08** (2013) 090 [[arXiv:1304.4926](#)] [[INSPIRE](#)].
- [6] X. Dong, A. Lewkowycz and M. Rangamani, *Deriving covariant holographic entanglement*, *JHEP* **11** (2016) 028 [[arXiv:1607.07506](#)] [[INSPIRE](#)].
- [7] B. Swingle, *Entanglement renormalization and holography*, *Phys. Rev. D* **86** (2012) 065007 [[arXiv:0905.1317](#)] [[INSPIRE](#)].
- [8] B. Swingle, *Constructing holographic spacetimes using entanglement renormalization*, [arXiv:1209.3304](#) [[INSPIRE](#)].
- [9] F. Pastawski, B. Yoshida, D. Harlow and J. Preskill, *Holographic quantum error-correcting codes: toy models for the bulk/boundary correspondence*, *JHEP* **06** (2015) 149 [[arXiv:1503.06237](#)] [[INSPIRE](#)].
- [10] P. Hayden, S. Nezami, X.-L. Qi, N. Thomas, M. Walter and Z. Yang, *Holographic duality from random tensor networks*, *JHEP* **11** (2016) 009 [[arXiv:1601.01694](#)] [[INSPIRE](#)].
- [11] A.J. Ferris and D. Poulin, *Tensor networks and quantum error correction*, *Phys. Rev. Lett.* **113** (2014) 030501 [[arXiv:1312.4578](#)].
- [12] T. Kohler and T. Cubitt, *Toy models of holographic duality between local Hamiltonians*, *JHEP* **08** (2019) 017 [[arXiv:1810.08992](#)] [[INSPIRE](#)].
- [13] A. Almheiri, X. Dong and D. Harlow, *Bulk locality and quantum error correction in AdS/CFT*, *JHEP* **04** (2015) 163 [[arXiv:1411.7041](#)] [[INSPIRE](#)].

- [14] Z. Yang, P. Hayden and X.-L. Qi, *Bidirectional holographic codes and sub-AdS locality*, *JHEP* **01** (2016) 175 [[arXiv:1510.03784](#)] [[INSPIRE](#)].
- [15] T. Takayanagi and K. Umemoto, *Entanglement of purification through holographic duality*, *Nature Phys.* **14** (2018) 573 [[arXiv:1708.09393](#)] [[INSPIRE](#)].
- [16] 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](#)].
- [17] X. Dong, *The gravity dual of Rényi entropy*, *Nature Commun.* **7** (2016) 12472 [[arXiv:1601.06788](#)] [[INSPIRE](#)].
- [18] P. Hayden, B. Swingle and M. Walter, *One-shot information theory in quantum field theory*, unpublished.
- [19] G. Vidal, *Entanglement renormalization*, *Phys. Rev. Lett.* **99** (2007) 220405 [[cond-mat/0512165](#)] [[INSPIRE](#)].
- [20] F. Verstraete, V. Murg and J. Cirac, *Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems*, *Adv. Phys.* **57** (2008) 143 [[arXiv:0907.2796](#)].
- [21] F. Verstraete and J.I. Cirac, *Renormalization algorithms for quantum-many body systems in two and higher dimensions*, [cond-mat/0407066](#).
- [22] T. Faulkner, A. Lewkowycz and J. Maldacena, *Quantum corrections to holographic entanglement entropy*, *JHEP* **11** (2013) 074 [[arXiv:1307.2892](#)] [[INSPIRE](#)].
- [23] N. Engelhardt and A.C. Wall, *Quantum extremal surfaces: holographic entanglement entropy beyond the classical regime*, *JHEP* **01** (2015) 073 [[arXiv:1408.3203](#)] [[INSPIRE](#)].
- [24] M. Freedman and M. Headrick, *Bit threads and holographic entanglement*, *Commun. Math. Phys.* **352** (2017) 407 [[arXiv:1604.00354](#)] [[INSPIRE](#)].
- [25] P. Hayden and A. Winter, *Communication cost of entanglement transformations*, *Phys. Rev. A* **67** (2003) 012326 [[quant-ph/0204092](#)].
- [26] B. Czech, P. Hayden, N. Lashkari and B. Swingle, *The information theoretic interpretation of the length of a curve*, *JHEP* **06** (2015) 157 [[arXiv:1410.1540](#)] [[INSPIRE](#)].
- [27] H. Wilming and J. Eisert, *Single-shot holographic compression from the area law*, *Phys. Rev. Lett.* **122** (2019) 190501 [[arXiv:1809.10156](#)] [[INSPIRE](#)].
- [28] M. Headrick, *Entanglement Rényi entropies in holographic theories*, *Phys. Rev. D* **82** (2010) 126010 [[arXiv:1006.0047](#)] [[INSPIRE](#)].
- [29] C. Akers and P. Rath, *Holographic Rényi entropy from quantum error correction*, *JHEP* **05** (2019) 052 [[arXiv:1811.05171](#)] [[INSPIRE](#)].
- [30] X. Dong, D. Harlow and D. Marolf, *Flat entanglement spectra in fixed-area states of quantum gravity*, *JHEP* **10** (2019) 240 [[arXiv:1811.05382](#)] [[INSPIRE](#)].
- [31] C. Fuchs and J. van de Graaf, *Cryptographic distinguishability measures for quantum-mechanical states*, *IEEE Trans. Inform. Theory* **45** (1999) 1216 [[quant-ph/9712042](#)].
- [32] G.C. Dorsch, S.J. Huber, T. Konstandin and J.M. No, *A second Higgs doublet in the early universe: baryogenesis and gravitational waves*, *JCAP* **05** (2017) 052 [[arXiv:1611.05874](#)] [[INSPIRE](#)].

- [33] D. Harlow, *The Ryu-Takayanagi formula from quantum error correction*, *Commun. Math. Phys.* **354** (2017) 865 [[arXiv:1607.03901](#)] [[INSPIRE](#)].
- [34] Y.Y. Shi, L.M. Duan and G. Vidal, *Classical simulation of quantum many-body systems with a tree tensor network*, *Phys. Rev. A* **74** (2006) 022320 [[quant-ph/0511070](#)] [[INSPIRE](#)].
- [35] L. Drescher and O. Fawzi, *On simultaneous min-entropy smoothing*, *Proc. ISIT* (2013) 161 [[arXiv:1312.7642](#)].
- [36] X. Dong, D. Harlow and A.C. Wall, *Reconstruction of bulk operators within the entanglement wedge in gauge-gravity duality*, *Phys. Rev. Lett.* **117** (2016) 021601 [[arXiv:1601.05416](#)] [[INSPIRE](#)].
- [37] J. Cotler, P. Hayden, G. Penington, G. Salton, B. Swingle and M. Walter, *Entanglement wedge reconstruction via universal recovery channels*, *Phys. Rev. X* **9** (2019) 031011 [[arXiv:1704.05839](#)] [[INSPIRE](#)].
- [38] P. Hayden and G. Penington, *Learning the alpha-bits of black holes*, [arXiv:1807.06041](#) [[INSPIRE](#)].
- [39] T. Faulkner and A. Lewkowycz, *Bulk locality from modular flow*, *JHEP* **07** (2017) 151 [[arXiv:1704.05464](#)] [[INSPIRE](#)].
- [40] P. Hayden and A. Winter, *Weak decoupling duality and quantum identification*, *IEEE Trans. Inform. Theory* **58** (2012) 4914 [[arXiv:1003.4994](#)].
- [41] P. Hayden and G. Penington, *Approximate quantum error correction revisited: introducing the alpha-bit*, [arXiv:1706.09434](#).
- [42] A. Hamilton, D.N. Kabat, G. Lifschytz and D.A. Lowe, *Holographic representation of local bulk operators*, *Phys. Rev. D* **74** (2006) 066009 [[hep-th/0606141](#)] [[INSPIRE](#)].
- [43] B.M. Terhal, M. Horodecki, D.W. Leung and D.P. Di Vincenzo, *The entanglement of purification*, *J. Math. Phys.* **43** (2002) 4286 [[quant-ph/0202044](#)].
- [44] G. Vidal, *On the characterization of entanglement*, *J. Mod. Opt.* **47** (2000) 355 [[quant-ph/9807077](#)] [[INSPIRE](#)].
- [45] N. Bao, *Minimal purifications, wormhole geometries and the complexity=action proposal*, [arXiv:1811.03113](#) [[INSPIRE](#)].
- [46] N. Bao and I.F. Halpern, *Holographic inequalities and entanglement of purification*, *JHEP* **03** (2018) 006 [[arXiv:1710.07643](#)] [[INSPIRE](#)].
- [47] N. Bao and I.F. Halpern, *Conditional and multipartite entanglements of purification and holography*, *Phys. Rev. D* **99** (2019) 046010 [[arXiv:1805.00476](#)] [[INSPIRE](#)].
- [48] K. Umemoto and Y. Zhou, *Entanglement of purification for multipartite states and its holographic dual*, *JHEP* **10** (2018) 152 [[arXiv:1805.02625](#)] [[INSPIRE](#)].
- [49] N. Bao, A. Chatwin-Davies and G.N. Remmen, *Entanglement of purification and multiboundary wormhole geometries*, *JHEP* **02** (2019) 110 [[arXiv:1811.01983](#)] [[INSPIRE](#)].
- [50] M. Miyaji and T. Takayanagi, *Surface/state correspondence as a generalized holography*, *PTEP* **2015** (2015) 073B03 [[arXiv:1503.03542](#)] [[INSPIRE](#)].
- [51] X.-L. Qi, Z. Yang and Y.-Z. You, *Holographic coherent states from random tensor networks*, *JHEP* **08** (2017) 060 [[arXiv:1703.06533](#)] [[INSPIRE](#)].

- [52] M. Han and S. Huang, *Discrete gravity on random tensor network and holographic Rényi entropy*, *JHEP* **11** (2017) 148 [[arXiv:1705.01964](#)] [[INSPIRE](#)].
- [53] A. Almheiri, X. Dong and B. Swingle, *Linearity of holographic entanglement entropy*, *JHEP* **02** (2017) 074 [[arXiv:1606.04537](#)] [[INSPIRE](#)].
- [54] N. Bao and H. Ooguri, *Distinguishability of black hole microstates*, *Phys. Rev. D* **96** (2017) 066017 [[arXiv:1705.07943](#)] [[INSPIRE](#)].
- [55] M. Bañados, C. Teitelboim and J. Zanelli, *Black hole entropy and the dimensional continuation of the Gauss-Bonnet theorem*, *Phys. Rev. Lett.* **72** (1994) 957 [[gr-qc/9309026](#)] [[INSPIRE](#)].
- [56] S. Carlip and C. Teitelboim, *Aspects of black hole quantum mechanics and thermodynamics in (2 + 1)-dimensions*, *Phys. Rev. D* **51** (1995) 622 [[gr-qc/9405070](#)] [[INSPIRE](#)].
- [57] D.L. Jafferis, A. Lewkowycz, J. Maldacena and S.J. Suh, *Relative entropy equals bulk relative entropy*, *JHEP* **06** (2016) 004 [[arXiv:1512.06431](#)] [[INSPIRE](#)].
- [58] N. Bao, C. Cao, S. Fischetti and C. Keeler, *Towards bulk metric reconstruction from extremal area variations*, *Class. Quant. Grav.* **36** (2019) 185002 [[arXiv:1904.04834](#)] [[INSPIRE](#)].
- [59] D. Marolf and A.C. Wall, *Eternal black holes and superselection in AdS/CFT*, *Class. Quant. Grav.* **30** (2013) 025001 [[arXiv:1210.3590](#)] [[INSPIRE](#)].
- [60] A. Almheiri, D. Marolf, J. Polchinski and J. Sully, *Black holes: complementarity or firewalls?*, *JHEP* **02** (2013) 062 [[arXiv:1207.3123](#)] [[INSPIRE](#)].
- [61] A. Almheiri, D. Marolf, J. Polchinski, D. Stanford and J. Sully, *An apology for firewalls*, *JHEP* **09** (2013) 018 [[arXiv:1304.6483](#)] [[INSPIRE](#)].
- [62] S.D. Mathur, *The information paradox: a pedagogical introduction*, *Class. Quant. Grav.* **26** (2009) 224001 [[arXiv:0909.1038](#)] [[INSPIRE](#)].
- [63] S.L. Braunstein, S. Pirandola and K. Życzkowski, *Better late than never: information retrieval from black holes*, *Phys. Rev. Lett.* **110** (2013) 101301 [[arXiv:0907.1190](#)] [[INSPIRE](#)].
- [64] K. Papadodimas and S. Raju, *State-dependent bulk-boundary maps and black hole complementarity*, *Phys. Rev. D* **89** (2014) 086010 [[arXiv:1310.6335](#)] [[INSPIRE](#)].
- [65] J. Maldacena and L. Susskind, *Cool horizons for entangled black holes*, *Fortsch. Phys.* **61** (2013) 781 [[arXiv:1306.0533](#)] [[INSPIRE](#)].
- [66] B. Czech, J.L. Karczmarek, F. Nogueira and M. Van Raamsdonk, *The gravity dual of a density matrix*, *Class. Quant. Grav.* **29** (2012) 155009 [[arXiv:1204.1330](#)] [[INSPIRE](#)].
- [67] V.E. Hubeny, *Extremal surfaces as bulk probes in AdS/CFT*, *JHEP* **07** (2012) 093 [[arXiv:1203.1044](#)] [[INSPIRE](#)].
- [68] N. Engelhardt and A.C. Wall, *Extremal surface barriers*, *JHEP* **03** (2014) 068 [[arXiv:1312.3699](#)] [[INSPIRE](#)].
- [69] V. Balasubramanian, B.D. Chowdhury, B. Czech and J. de Boer, *Entwinement and the emergence of spacetime*, *JHEP* **01** (2015) 048 [[arXiv:1406.5859](#)] [[INSPIRE](#)].
- [70] Y. Nomura, P. Rath and N. Salzetta, *Pulling the boundary into the bulk*, *Phys. Rev. D* **98** (2018) 026010 [[arXiv:1805.00523](#)] [[INSPIRE](#)].

- [71] V. Balasubramanian, P. Hayden, A. Maloney, D. Marolf and S.F. Ross, *Multiboundary wormholes and holographic entanglement*, *Class. Quant. Grav.* **31** (2014) 185015 [[arXiv:1406.2663](#)] [[INSPIRE](#)].
- [72] D. Stanford and L. Susskind, *Complexity and shock wave geometries*, *Phys. Rev. D* **90** (2014) 126007 [[arXiv:1406.2678](#)] [[INSPIRE](#)].
- [73] A. Peach and S.F. Ross, *Tensor network models of multiboundary wormholes*, *Class. Quant. Grav.* **34** (2017) 105011 [[arXiv:1702.05984](#)] [[INSPIRE](#)].
- [74] P. Hayden, M. Headrick and A. Maloney, *Holographic mutual information is monogamous*, *Phys. Rev. D* **87** (2013) 046003 [[arXiv:1107.2940](#)] [[INSPIRE](#)].
- [75] N. Bao, S. Nezami, H. Ooguri, B. Stoica, J. Sully and M. Walter, *The holographic entropy cone*, *JHEP* **09** (2015) 130 [[arXiv:1505.07839](#)] [[INSPIRE](#)].
- [76] S.X. Cui, P. Hayden, T. He, M. Headrick, B. Stoica and M. Walter, *Bit threads and holographic monogamy*, [arXiv:1808.05234](#) [[INSPIRE](#)].
- [77] V.E. Hubeny, *Bulk locality and cooperative flows*, *JHEP* **12** (2018) 068 [[arXiv:1808.05313](#)] [[INSPIRE](#)].
- [78] C.A. Agón, J. De Boer and J.F. Pedraza, *Geometric aspects of holographic bit threads*, *JHEP* **05** (2019) 075 [[arXiv:1811.08879](#)] [[INSPIRE](#)].
- [79] O. Aharony, O. Bergman, D.L. Jafferis and J. Maldacena, *$N = 6$ superconformal Chern-Simons-matter theories, M2-branes and their gravity duals*, *JHEP* **10** (2008) 091 [[arXiv:0806.1218](#)] [[INSPIRE](#)].
- [80] A. Mollabashi, N. Shiba and T. Takayanagi, *Entanglement between two interacting CFTs and generalized holographic entanglement entropy*, *JHEP* **04** (2014) 185 [[arXiv:1403.1393](#)] [[INSPIRE](#)].
- [81] A. Karch and C.F. Uhlemann, *Holographic entanglement entropy and the internal space*, *Phys. Rev. D* **91** (2015) 086005 [[arXiv:1501.00003](#)] [[INSPIRE](#)].
- [82] X.-L. Qi and Z. Yang, *Space-time random tensor networks and holographic duality*, [arXiv:1801.05289](#) [[INSPIRE](#)].
- [83] A. May, *Tensor networks for dynamic spacetimes*, *JHEP* **06** (2017) 118 [[arXiv:1611.06220](#)] [[INSPIRE](#)].
- [84] A.C. Wall, *Maximin surfaces and the strong subadditivity of the covariant holographic entanglement entropy*, *Class. Quant. Grav.* **31** (2014) 225007 [[arXiv:1211.3494](#)] [[INSPIRE](#)].
- [85] Y. Chen, X. Dong, A. Lewkowycz and X.-L. Qi, *Modular flow as a disentangler*, *JHEP* **12** (2018) 083 [[arXiv:1806.09622](#)] [[INSPIRE](#)].
- [86] T. Faulkner, M. Li and H. Wang, *A modular toolkit for bulk reconstruction*, *JHEP* **04** (2019) 119 [[arXiv:1806.10560](#)] [[INSPIRE](#)].
- [87] R. Bousso, *Holography in general space-times*, *JHEP* **06** (1999) 028 [[hep-th/9906022](#)] [[INSPIRE](#)].
- [88] F. Sanches and S.J. Weinberg, *Holographic entanglement entropy conjecture for general spacetimes*, *Phys. Rev. D* **94** (2016) 084034 [[arXiv:1603.05250](#)] [[INSPIRE](#)].