

Discrete bulk reconstruction

Scott Aaronson and Jason Pollack

*Department of Computer Science, University of Texas at Austin,
2317 Speedway, Austin, TX, U.S.A.*

E-mail: scott@scottaaronson.com, jasonpollack@gmail.com

ABSTRACT: According to the *AdS/CFT correspondence*, the geometries of certain space-times are fully determined by quantum states that live on their boundaries — indeed, by the von Neumann entropies of portions of those boundary states. This work investigates to what extent the geometries can be reconstructed from the entropies *in polynomial time*. Bouland, Fefferman, and Vazirani (2019) argued that the AdS/CFT map can be exponentially complex if one wants to reconstruct regions such as the interiors of black holes. Our main result provides a sort of converse: we show that, in the special case of a single 1D boundary divided into N “atomic regions”, if the input data consists of a list of entropies of *contiguous* boundary regions, and if the entropies satisfy a single inequality called Strong Subadditivity, then we can construct a graph model for the bulk in linear time. Moreover, the bulk graph is planar, it has $O(N^2)$ vertices (the information-theoretic minimum), and it’s “universal”, with only the edge weights depending on the specific entropies in question. From a combinatorial perspective, our problem boils down to an “inverse” of the famous min-cut problem: rather than being given a graph and asked to find a min-cut, here we’re given the values of min-cuts separating various sets of vertices, and need to find a weighted undirected graph consistent with those values. Our solution to this problem relies on the notion of a “bulkless” graph, which might be of independent interest for AdS/CFT. We also make initial progress on the case of multiple 1D boundaries — where the boundaries could be connected via wormholes — including an upper bound of $O(N^4)$ vertices whenever an embeddable bulk graph exists (thus putting the problem into the complexity class NP).

KEYWORDS: AdS-CFT Correspondence, Gauge-Gravity Correspondence

ARXIV EPRINT: [2210.15601](https://arxiv.org/abs/2210.15601)

Contents

1	Introduction	1
1.1	Our results	3
1.2	The Ryu-Takayanagi formula and min-cut	4
1.3	Previous work	5
2	Preliminaries	5
2.1	Statement of the problem	5
2.2	Computability with doubly exponential vertices	7
2.3	Planarity	8
2.4	Entropic inequalities	9
2.5	Entropies from contiguous data	11
3	The single 1D boundary case	13
3.1	Bulkless graphs	13
3.2	The chord construction	16
3.3	The diamondwork construction	17
3.4	Diamondwork examples	20
4	The multiple-boundary case	23
4.1	2-boundary counterexamples	24
4.2	Size upper bound for embeddable graphs	28
5	Discussion and open problems	30

1 Introduction

The *anti-de Sitter/conformal field theory* or AdS/CFT correspondence is one of the most important developments in theoretical physics of the past quarter century. AdS/CFT posits a duality between two superficially different theories:

- (i) a theory of quantum gravity in a D -dimensional “bulk” space with a negative cosmological constant (called anti-de Sitter or AdS space), and
- (ii) a quantum field theory, with no gravity, that lives on¹ the $(D-1)$ -dimensional boundary of the D -dimensional AdS space.

¹Strictly speaking, since the two theories share the same Hilbert space, the field theory does not literally live on the boundary of the bulk space, but rather on a *copy* of it: that is, on a space isometric to the boundary. We will often elide such distinctions when speaking informally.




Figure 1. Representation of the hyperbolic plane on the Poincaré disk. White lines denote boundary-anchored geodesics. The piling up of fish at the boundary indicates that lengths become infinite as they approach it. (Figure from <https://www.d.umn.edu/~ddunham/dunbrid07.pdf>, modeled on *Circle Limit III*, M.C. Escher, 1959.)

The relation between the two theories is often called “holographic”; it necessarily fails to respect spatial locality. Though it originally emerged from string theory, AdS/CFT is now often discussed in the context of the black hole information problem and so on with no explicit reference to strings.

The universe described by AdS/CFT is not our universe: among other differences, it has a negative cosmological constant whereas ours appears to have a positive one. Nevertheless, AdS/CFT stands as the most explicit known example where a classical spacetime emerges from more fundamental, quantum degrees of freedom, as envisioned by the “It from Qubit” approach to physics.

What does it mean for the bulk and boundary theories to be “equivalent”? It means that there’s a mapping, or “dictionary”, mapping all states and observables in one theory to corresponding states and observables in the other, justifying the view that the two theories are just different ways to look at the same underlying Hilbert space. This dictionary is not yet completely known² — indeed, the AdS side lacks an independent fully-rigorous definition — but enough calculations have been done in both theories and have yielded the same answers to convince most experts that the dictionary exists.

A key entry in the holographic dictionary is the *Ryu-Takayanagi (RT) formula* [1], which equates the von Neumann entropy of a subregion of the boundary state with the minimal area among all bulk surfaces that end on that boundary subregion (see figure 1). Hence, if we are given a CFT state, we could imagine systematically *reconstructing* the bulk geometry by looking at the entropies of various boundary subregions to determine

²Some string theorists would disagree.

areas. In practice, we will not be able to reconstruct beyond event horizons in the bulk spacetime (or, more generally, extremal surface barriers [2]), but we can nevertheless often learn the metric in a large portion of the spacetime [3, 4].

The RT formula is extremely interesting from a quantum-information perspective because it implies that, in holographic states, entropies behave like areas, and thus must satisfy all the same constraints that areas do. Entropies in *all* quantum states satisfy certain inequalities, such as *subadditivity* and *strong subadditivity*:

$$S(\rho_{AB}) \leq S(\rho_A) + S(\rho_B), \quad S(\rho_{ABC}) + S(\rho_B) \leq S(\rho_{AB}) + S(\rho_{BC}). \quad (1.1)$$

But entropies in *holographic* states must satisfy additional inequalities, such as *monogamy of mutual information* [5]:

$$S(\rho_A) + S(\rho_B) + S(\rho_C) + S(\rho_{ABC}) \leq S(\rho_{AB}) + S(\rho_{BC}) + S(\rho_{AC}), \quad (1.2)$$

as well as an infinite family of further inequalities [6]. Note that the difficult-sounding problem of which boundary quantum states describe classical bulk geometries thereby gets reduced to a much more concrete problem, of checking that a list of entropy inequalities is satisfied.

It is natural to wonder just how simple AdS/CFT can be made. Can we see the main ideas in a model scenario, without using the machinery of quantum field theory? Work over the past decade has aimed to answer this: for instance, [7, 8] interpreted AdS/CFT as literally an example of a quantum error-correcting code, while [9, 10] proposed tensor networks which exhibit a form of the RT formula. Meanwhile, [6] proposed modeling the bulk space by a weighted graph with finitely many vertices. In this case, minimal surfaces turn into *min-cuts*: that is, sets of edges of minimum total weight that separate the graph into disconnected components.

In this work, we seek to push the simplification process as far as possible. We dispense entirely with the bulk spatial manifold and the boundary quantum state, and just proceed directly from a list of entropies, for unions of “atomic” boundary regions, to a graph model of the bulk — i.e., a weighted graph whose min-cuts correspond to the given entropies. We ask both when this graph exists, and how easily it can be constructed when it does.

The mapping from entropies to a graph will not be unique. The mapping from the graph to a full bulk geometry will *also* not be unique. Nevertheless, the graph can be thought of as a particular coarse-graining of a geometry, or as the common data shared by an equivalence class of possible geometries.

1.1 Our results

In the important case $D = 2$, where the spatial boundary consists of one or more circles, we make significant progress. When the spatial boundary is a *single* circle divided into N “atomic regions”, we show that a planar bulk graph is fully determined by the entropies of the quadratically many *contiguous* boundary regions, and always exists provided those entropies satisfy Strong Subadditivity. Furthermore, this graph has $O(N^2)$ vertices and is *universal*, meaning that we can reproduce *any* valid vector of contiguous entropies merely

by adjusting the graph’s edge weights. The mapping from contiguous entropies to edge weights is just a linear transformation, and is computable in $O(N^2)$ time — linear in the amount of input data.

The fact that we’re only trying to explain the entropies of *contiguous* boundary regions is crucial here. Note that, *assuming* the bulk graph is planar, it’s already known how to pass from the vector of contiguous entropies to the entropy of any desired non-contiguous region: the work of [11] reduced that problem to an instance of minimum-weight bipartite perfect matching, which is solvable in polynomial time. On the other hand, there are also valid holographic states with non-contiguous entropies that violate the Bao-Chatwin-Davies prescription. We can interpret these states as describing bulk graphs that fail to be planar because they contain *wormholes* connecting faraway boundary regions.

Note also that, with $D \geq 3$ bulk dimensions, there are already $\exp(N)$ *contiguous* boundary regions, and the bulk reconstruction problem is much more complicated for that reason, among others.

When there are multiple boundary regions, the resulting bulk geometry may have wormholes connecting the boundaries, and the problem again becomes more complicated. Although our main graph construction is unable to handle the multi-boundary case, we prove some weaker results and no-go theorems.

1.2 The Ryu-Takayanagi formula and min-cut

For computer scientists, a striking feature of AdS/CFT is the role played by *min-cuts*, a central concept in combinatorics and graph theory. Given a finite weighted undirected graph G with real edge weights $w(e) \geq 0$, as well as two disjoint sets of vertices R, R' , a min-cut is a set C of edges of minimum total weight, $W = \sum_{e \in C} w(e)$, whose removal disconnects R from R' . The famous Max-Flow/Min-Cut theorem [12] says that W , the weight of the min-cut, equals the maximum amount of “flow” that can be routed from R to R' via the edges of G , where the flow along any edge cannot exceed its weight, and the total flow entering any vertex must equal the flow leaving it. Polynomial-time algorithms to compute the min-cut, or equivalently max-flow, between R and R' are staples of the undergraduate CS curriculum, e.g., [13]. Nearly linear-time algorithms for these problems were even recently achieved [14].

The connection between min-cuts and AdS/CFT comes via the RT formula discussed previously. Recall what the RT formula says: that, up to tiny corrections, the area of a *minimal surface* in the bulk separating a boundary region R from its complement is proportional to $S(\rho_R)$, the von Neumann entropy of the reduced quantum state of CFT on region R , with the coefficient of proportionality given by $1/(4G_N)$, with G_N the gravitational constant. When we replace manifolds by graphs, the geometric concept of a minimal surface becomes *precisely* the combinatorial concept of a min-cut. Likewise, the problem of constructing a bulk geometry to represent the boundary entropies, becomes the problem of constructing a graph with prescribed min-cut values between various distinguished sets of vertices.

Our concern, then, is with the “inverse” of the classic min-cut problem: rather than being given a graph and asked to find a min-cut, we will be given a list of min-cut values and asked to find a graph consistent with them.

1.3 Previous work

Beyond the “It from Qubit” program in general, this paper was specifically inspired by research into the entropies of subregions of holographic states — a line of work initiated by [5, 6] and continued by, for example, [15–19].

For the most part, authors have focused either on generalizing to classes of states beyond holographic states, or better mathematical understanding of the structure of the space of holographic states, or ways of obtaining the Ryu-Takayanagi formula using tools other than minimal surfaces. Our work does not fall into any of these categories. Instead, we work directly with minimal areas, and are concerned with reproducing them as the min-cuts of a graph.

As we review in section 2.2, [6] already considered the problem of going back and forth between geometries and graphs, and gave a general method to obtain a graph given a geometry and a set of subregions. Unfortunately, their method generically produces a graph with a number of vertices *doubly exponential* in the number of regions. [11] later explicitly studied the complexity of obtaining the minimal surfaces of regions in a given geometry.

One central point of departure from these earlier works is that we never assume a manifold or a metric. Instead, we take as input only a finite list of the entropies themselves, and then try to pass *directly* to a graph model of the bulk — now taking care to minimize the graph’s size, as well as the computational complexity of building the graph. As we’ll see, in an important special case — namely, when we’re given as input the entropies of contiguous regions along a 1D boundary — we’ll be able to construct a graph with $O(N^2)$ vertices in linear time, a vast improvement over the doubly-exponential generic construction of [6].

Our work can be seen as a sort of converse to the widely-discussed paper of [20], on the computational complexity of the holographic dictionary. Those authors argued that, when the bulk geometry contains black holes or wormholes, the problem of reconstructing the geometry from boundary data can be exponentially hard even for a quantum computer. Indeed, they showed this task to be at least as hard as distinguishing various candidate pseudorandom n -qubit states from Haar-random states. By contrast, this work aims to show that, when the bulk *lacks* event horizons or nontrivial topology (or, more generally, the extremal surfaces which create the “python’s lunch” phenomenon [21, 22]), its geometry can be completely reconstructed from boundary data in classical polynomial time. We make significant progress toward this conjecture by proving it in the special case of a single 1D boundary.

2 Preliminaries

2.1 Statement of the problem

In the previous section, we introduced the problem of reconstructing a bulk graph from the entropies of boundary regions. In general, we might only have access to the entropies of a limited set of subregions. Given a factorization of a finite-dimensional Hilbert space

as an N -fold tensor-product,

$$\mathcal{H} = \otimes_{i=1}^N \mathcal{H}_i, \quad (2.1)$$

we can group the entropies of the various reduced density matrices constructible from a state $|\Psi\rangle \in \mathcal{H}$ into an *entropy vector* of length $2^{N-1} - 1$,

$$\vec{S}(|\Psi\rangle) \equiv (S_1, S_2, \dots, S_N, S_{1,2}, S_{1,3}, \dots, S_{1,N}, \dots, S_{N-1,N}, S_{1,2,3}, \dots), \quad (2.2)$$

where $S_{a_1, \dots, a_k}$ is the von Neumann entropy of the reduced state $\rho_{a_1, \dots, a_k}$ on $a_1, \dots, a_k$, and the indices are over all sets of elements in $\{1, \dots, N\}$, except that because of the purity of $|\Psi\rangle$ there is a redundancy between the entropy of a set and that of its complement (and $S_\emptyset = S_{1, \dots, N} = 0$ is omitted).

One possible approach would therefore be to start explicitly with a quantum state $|\Psi\rangle \in \mathcal{H}$ and work with its entropy vector relative to some factorization. However, specifying the state in, say, the field value basis would require a continuous function's worth of degrees of freedom, which would defeat our goal of working with purely discrete objects. Indeed, the Hilbert space of a conformal field theory is infinite-dimensional and not isomorphic to the tensor product of a finite number of smaller factors. In principle, these problems could be cured by explicitly fixing an ultraviolet and infrared cutoff for the field theory, so that the Hilbert space became explicitly finite-dimensional, but in practice this procedure is difficult to carry out and sensitive to details of how the cutoff is implemented.

Our approach is simpler: we drop the demand that the entropies be derived from some particular quantum state $|\Psi\rangle$. Instead, we take our input to be an abstract list of entropies, either $2^{N-1} - 1$ entries arranged into an entropy vector,

$$\vec{S}^{(N)} \equiv (S_1, S_2, \dots, S_N, S_{1,2}, S_{1,3}, \dots, S_{1,N}, \dots, S_{N-1,N}, S_{1,2,3}, \dots), \quad (2.3)$$

or some subset of this list. For all $i \in [N]$, we will say that S_i is an entropy of the i^{th} *atomic region*, and have in mind that the union of the N atomic regions comprises the entire boundary. For example, as shown in figure 2, we could think of the N atomic regions as subintervals of a spatial circle. In this case, it would be natural to think of the entropy vector $\vec{S}^{(N)}$ as derived from some collection of qubits (or qudits) living on the circle, but we will not actually demand this. (We will return in section 5 to the question of when we can construct a concrete quantum state with a specified entropy vector.)

To avoid excessive notation, we often label atomic regions by Latin letters $A, B, C, \dots$. We write, for example, $S(AC)$ for the entropy of $AC \equiv A \cup C$, the union of the atomic regions A and C .

The discrete bulk reconstruction problem can now be formulated as follows:

Problem 1 (discrete bulk reconstruction problem (DBRP)). *We're given as input a list of atomic boundary regions labeled $1, \dots, N$, a list of subsets of the regions $R_1, \dots, R_k \subseteq [N]$, and a real-valued entropy $S(R_i) \geq 0$ for each R_i . The problem is to construct a weighted undirected graph G , with N distinguished boundary vertices $1, \dots, N$ that we identify with the atomic boundary regions, such that for each R_i , the weight of the minimum cut separating R_i from the rest of the boundary vertices (i.e., from $[N] - R_i$) is equal to $S(R_i)$. (Or to output that no such graph exists.)*




Figure 2. Division of the boundary into atomic regions.

When such a graph does exist, we would like to know if it is essentially unique, how many vertices it has, whether it is planar, and so on.

2.2 Computability with doubly exponential vertices

A priori, one might worry that the DBRP could be uncomputable (equivalent to the halting problem), with no upper bound $f(N)$ on the number of bulk vertices needed as a function of the number N of boundary vertices. Fortunately, the work of [6] showed that this is not the case, and that $f(N) = 2^{2^N}$ vertices always suffice. Let us prove this for completeness.

Proposition 2 ([6]). *Whenever a graph exists that solves the DBRP on a given set of input data, there exists a (possibly identical) graph with at most 2^{2^N} vertices that also solves it.*

Proof. By definition, any graph G solving the DBRP has N boundary vertices $\{v_1, \dots, v_N\}$. For every subset $S \subset \{v_1, \dots, v_N\}$ of the boundary vertices, find the min-cut separating S from $\{v_1, \dots, v_N\} - S$. The min-cut divides the graph into two subgraphs; call the one containing S the “RT region of S .” There are 2^N distinct subsets of boundary vertices, and hence at most 2^N RT regions. Then construct all 2^{2^N} intersections of the RT regions (some of which may be the empty set). If there exists an intersection $H \subseteq G$ with more than one vertex, then we can construct a strictly smaller graph G' that also solves the DBRP with the same input data, by preserving all edges which cross the boundary of H , and connecting them all to a single internal vertex. See figure 3. Repeat this procedure until no intersections with more than one vertex remain. Then the total number of vertices in the graph is at most 2^{2^N} . $\square$

Proposition 2 has the following important consequence.

Corollary 3. *The DBRP is Turing-computable.*




Figure 3. A minimal graph with one vertex for every intersection of RT regions. (For a non-planar graph there will in general be many more non-empty intersections, not shown here.)

Proof. To solve the DBRP for given input data, we first construct the graph of size 2^{2^N} from proposition 2, with one vertex for each intersection. We then write down a list of $\exp(2^{2^N})$ inequalities on the edge weights to express that each of the 2^N min-cuts have at least the correct values. Next, we loop over all $\exp(2^{2^N})^{2^N}$ tuples of cuts to be forced to have at *most* the correct values. Finally, for each such tuple, we search for a solution to the resulting system of inequalities using linear programming, halting whenever a solution is found. This takes $\exp(\exp(N))$ time overall. $\square$

Of course, we would like to be more efficient whenever possible.

2.3 Planarity

When trying to construct a graph representing a 2D bulk, one extremely natural requirement is that the graph be planar. A planar graph seems interpretable as a discrete analogue of a 2-manifold with trivial topology (so in particular, no wormholes). One of the main surprises of this paper is that, given any entropies for contiguous boundary regions, we can always construct a *planar* bulk graph.

Having said that, we'll see that planarity can break down when there are entropies of *non-contiguous* boundary regions to account for — an especially natural situation when there are multiple 1D boundaries. In this case, we can effectively force there to be wormholes in the bulk.

To show this, we'll use the famous Kuratowski's theorem [23], which says that a graph G is planar if and only if it does not contain either the complete graph K_5 or the complete bipartite graph $K_{3,3}$ as minors (i.e., if neither can be obtained from G via contractions and deletions). Thus, it suffices to construct boundary regions whose entropies can only be explained via K_5 or $K_{3,3}$. The set of geodesics which start and end on these boundary regions can thus not be embedded into a 2-manifold; the geometry of a holographic state

of which such a graph is a discretization must therefore contain wormholes through which some of these geodesics pass. (Heuristically, we'd expect a graph embedding $K_{3,3}$ to be embedded in a graph with a multi-throated wormhole, while a graph containing K_5 might be embedded in a geometry with multiple separate two-sided wormholes, but we have not investigated this question in detail.)

2.4 Entropic inequalities

Among vectors of $2^{N-1} - 1$ nonnegative real numbers, not all are valid lists of entropies for the subsystems of some N -partite pure quantum state. The set that's valid is called the *entropy cone*: “cone” because it can be shown to be closed under nonnegative linear combinations (for a review, see e.g., [24]). Two inequalities which are valid for any quantum state are subadditivity (SA),

$$S(\rho_{AB}) \leq S(\rho_A) + S(\rho_B), \quad (2.4)$$

and strong subadditivity (SSA),

$$S(\rho_{ABC}) + S(\rho_B) \leq S(\rho_{AB}) + S(\rho_{BC}), \quad (2.5)$$

where A , B , and C can be any 3 disjoint subsystems. (Strong subadditivity implies subadditivity, as can be seen by taking the B system to be trivial, but it will sometimes be convenient to refer to the latter separately.) For larger N , the complete list of inequalities satisfied by the entropies of arbitrary N -partite quantum states is not yet known, nor is it even known to be finite.

Within the entropy cone is the so-called *holographic entropy cone*: the set of entropy vectors that correspond to N -partite quantum states with holographic duals. Formally, [6] defined the N^{th} holographic entropy cone as simply the set of entropy vectors $v \in \mathbb{R}^{2^N-1}$ for which there exists a weighted, undirected graph G , with N boundary vertices and any finite number of bulk vertices, such that for all subsets $R \subseteq [N]$ of the boundary vertices, the value of a min-cut in G separating R from $[N] - R$ is exactly $S(R)$. They justified this definition by arguing that a holographic bulk dual, in the physics sense, can always be converted to such a graph and vice versa.

It's known that, for every N , the holographic entropy cone is defined by a finite list of inequalities [6], and those inequalities can in principle be computed in finite time given N . It's known further that, for a pure state divided into $N \leq 3$ regions, all true entropy inequalities follow from SSA, meaning that the entropy cone and the holographic entropy cone coincide. For $N \geq 4$, however, there are additional inequalities that cause the holographic entropy cone to be a strict subset of the full entropy cone [25].

One such inequality is called *monogamy of mutual information*, or MMI:

$$S(\rho_{AB}) + S(\rho_{BC}) + S(\rho_{AC}) \geq S(\rho_A) + S(\rho_B) + S(\rho_C) + S(\rho_{ABC}). \quad (2.6)$$

All known inequalities that define the holographic entropy cone have proofs that consist of “cutting and pasting” geodesics. As an example, we show cutting-and-pasting proofs of subadditivity, strong subadditivity, and MMI in figures 4–6. We raise, as an open question, whether there are inequalities defining the holographic entropy cone that don't have cutting-and-pasting proofs of this kind.




Figure 4. Cutting and pasting proof of Subadditivity. To form a curve (though not necessarily a minimal one) that separates the region AB from its complement, we can simply take the union of a curve for A and a curve for B , as in the triangle inequality.




Figure 5. Cutting and pasting proof of Strong Subadditivity. The union of the blue and red curves has length $S(AB) + S(BC)$. But the blue curve ends on ABC , and thus has length at least $S(ABC)$, while the red curve ends on B , and thus has length at least $S(B)$.




Figure 6. Cutting and pasting proof of MMI. The proof is divided into two cases according to whether $S(AC)$ (in purple) is equal to $S(A) + S(C)$ (case (1)) or $S(ABC) + S(B)$ (case (2)); the two cases correspond to the disconnected and connected cases in figure 7, respectively. In both cases, the union of the blue and red curves has length $S(AB) + S(BC)$. In case (1), the length of the red curve is lower-bounded by $S(ABC)$ and the length of the blue curve is lower-bounded by $S(B)$. In case (2), the length of the red curve is lower-bounded by $S(A)$ and the length of the blue curve is lower-bounded by $S(C)$. (The holographic proof of MMI in [5] combines these cases, but we have separated them here for clarity.)

2.5 Entropies from contiguous data

Given an ordering of N atomic regions on a single boundary, we can easily check whether a set of regions is *contiguous*, that is, whether their union consists of only a single connected component. For example, AB and BC are contiguous but (for $N > 3$) AC is not. Furthermore, the complement of a contiguous region is itself contiguous.

Hence there are $N(N-1)/2 = \binom{N}{2}$ independent entropies of contiguous regions. (When N is even, we only need half of the size- $N/2$ regions.)

One might wonder: what is so special about the *contiguous* boundary regions, that could justify finding a bulk model to explain their entropies alone? The significance is that, as observed for example by [11]:

Proposition 4 (contiguous data suffices). *If there is a single 1D boundary, and if the bulk geometry is topologically trivial (e.g., does not contain wormholes), then the entropies of the contiguous boundary regions determine the entropies of the non-contiguous boundary regions as well.*

The proof of proposition 4 is simply that, in the situation described, the RT surface of a non-contiguous region R must be the union of RT surfaces of contiguous regions, or equivalently, geodesics that start and end on the boundary. (This is so because, if some component of R 's RT surface were *not* a geodesic, then we could decrease the total area by replacing it with one.) But knowing the entropies of all the contiguous regions tells us the lengths of all boundary-anchored geodesics.

Thus, suppose R is the union of k contiguous boundary regions $R_1, \dots, R_k$. Then to calculate the entropy $S(R)$, we “merely” need to solve a combinatorial optimization problem: namely, to take the minimum, over all sets of geodesics $g_1, \dots, g_k$ whose union separates R from its complement, of $l(g_1) + \dots + l(g_k)$, where $l(g_i)$ is the length of g_i . To illustrate, in figure 7, we can calculate the entropy of the non-contiguous region AC as

$$S(AC) = \min \{S(A) + S(C), S(B) + S(D)\}, \quad (2.7)$$

minimizing over the two possibilities for how to separate AC from BD .

We remark that we implicitly assumed the validity of proposition 4, and (2.7) in particular, in our proof of MMI in figure 6. One way to violate (2.7) is by adding an EPR pair between A and C , producing a state that violates MMI and is therefore non-holographic. Another is to change the topology of the boundary, e.g. by adding a second boundary circle connected in the bulk by a wormhole, allowing for other minimal surfaces beyond those possible for a single 1D boundary and trivial bulk topology.

In general, with k contiguous boundary regions $R_1, \dots, R_k$, there are $k!$ possibilities for geodesics separating $R = R_1 \cup \dots \cup R_k$ from its complement. As observed by [11], the problem of minimizing over those $k!$ possibilities can be cast as an instance of the minimum-weight perfect matching problem, for a bipartite graph H with k vertices on each side. The left vertices of this H correspond to the left endpoints of the R_i 's as we proceed around the circle clockwise, the right vertices correspond to their right endpoints,




Figure 7. Obtaining the RT surfaces of disconnected regions as unions of RT surfaces of connected regions. The red curve is the union of the RT surface of A (with length $S(A)$) and the RT surface of C (with length $S(C)$); the blue curve is the union of the RT surface of $ABC \equiv D$ (with length $S(D)$) and the RT surface of B (with length $S(B)$). Both are extremal surfaces ending on $AC \equiv A \cup C$; we must therefore have $S(AC) \leq \min(S(A) + S(C), S(B) + S(D))$. For a single disconnected boundary, the validity of the RT formula implies that no further extremal surfaces exist, and the inequality is saturated. When the red curve is the RT surface, we say that the entanglement wedge of AC is disconnected; when the blue curve is the RT surface, we say it is connected.

and the weight of an edge (v, w) is just the length of the minimal geodesic connecting v and w . This matching problem is well-known to be solvable in time polynomial in k .

Hence, under the assumptions of a single 1D boundary and trivial topology, we can represent the areas of *exponentially many* RT surfaces — those of the non-contiguous regions — in terms of the areas of the $\binom{N}{2}$ RT surfaces of the contiguous regions. This fact is so useful that it motivates a definition: call an entropy vector $v \in \mathbb{R}_{\geq 0}^{2^{N-1}-1}$ a *matching vector* if all non-contiguous entropies in it are obtained from the $\binom{N}{2}$ contiguous entropies via minimization over perfect matchings, as in the prescription above. A matching vector, despite its exponential length, is fully determined by polynomially many parameters.

Here we should pause to discuss a subtlety. As explained in section 2.4, the set of valid holographic entropy vectors $v \in \mathbb{R}_{\geq 0}^{2^{N-1}-1}$ forms a *cone*; that is, it's closed under nonnegative linear combinations. The set of valid contiguous entropy vectors, $w \in \mathbb{R}_{\geq 0}^{\binom{N}{2}}$, *also* forms a cone. By contrast, the set of matching vectors — that is, the subset of $\mathbb{R}_{\geq 0}^{2^{N-1}-1}$ that's obtainable from some contiguous entropy vector $w \in \mathbb{R}_{\geq 0}^{\binom{N}{2}}$ via the matching prescription — *does not* form a cone.

Here is an example that shows this: let A, B, C, D be the $N = 4$ atomic boundary regions. In state ρ , we'll have

$$\begin{aligned} S(A) &= S(C) = 2, \\ S(B) &= S(D) = 1, \\ S(AC) &= S(BD) = \min\{S(A) + S(C), S(B) + S(D)\} = 2, \end{aligned} \tag{2.8}$$

with $S(AB) = S(CD)$ and $S(AD) = S(BC)$ taking any values consistent with Strong Subadditivity. In state σ , we'll have

$$\begin{aligned} S(A) &= S(C) = 1, \\ S(B) &= S(D) = 2, \\ S(AC) &= S(BD) = \min\{S(A) + S(C), S(B) + S(D)\} = 2. \end{aligned} \tag{2.9}$$

Hence in state $\rho + \sigma$, assuming convexity we must have

$$\begin{aligned} S(A) &= S(C) = 3, \\ S(B) &= S(D) = 3, \\ S(AC) &= S(BD) = 4. \end{aligned} \tag{2.10}$$

But

$$\min\{S(A) + S(C), S(B) + S(D)\} = 6 > 4,$$

which gives us our contradiction.

3 The single 1D boundary case

Having said what we could about the general DBRP in the previous section, in this section we radically specialize to the case where the N atomic regions live on a single, 1-dimensional boundary (homologous to a circle). Furthermore, we assume that the input data consists of entropies for *contiguous* boundary regions only; section 2.5 explained why this is a reasonable choice.

3.1 Bulkless graphs

Suppose we are given the $\binom{N}{2}$ independent contiguous entropies for N atomic regions on a single boundary. By construction, any graph that solves the DBRP for this data must have at least N vertices: the boundary vertices themselves. Assuming we don't care about planarity, how many additional vertices must be provided?

In this section, we show that the surprising answer is: zero additional vertices! More precisely, consider *any* vector $v \in \mathbb{R}_{\geq 0}^{\binom{N}{2}}$ of contiguous entropies that obeys strong subadditivity (2.5) — or, equivalently, that gives rise to a point in the quantum entropy cone. We show that v admits an N -vertex graph, which we call a *bulkless graph*, that solves the DBRP. The bulkless graph for $N = 7$ is shown in figure 8.

Lemma 5 (bulkless graphs). *Let a bulkless graph be a weighted undirected graph with only N boundary vertices $A, B, C, \dots$. Suppose we are given the $\binom{N}{2}$ entropies for the contiguous boundary regions, and suppose they obey strong subadditivity (SSA). Then there exists a bulkless graph (indeed, a unique such graph) that solves the DBRP for this data. The edge weights for this graph can be computed in $O(N^2)$ time (i.e., linear in the amount of input data).*




Figure 8. The bulkless graph for $N = 7$ vertices. We can think of the vertices as ordered around a circle and connected by all possible chords. Cuts are given by chords that separate the boundary vertices in a contiguous region R from those in its complement $[N] - R$: for example, the blue chord is a cut between ABC and $DEFG$.

Proof. A bulkless graph is just a complete graph on N vertices with appropriate edge weights. Our problem is how to assign nonnegative weights so that each min-cut of the graph gives the correct entropy.

Label the weight of an edge by the two boundary vertices it connects, e.g. w_{AB} is the weight of the edge connecting A and B . For any contiguous set of vertices L , there is a unique min-cut separating L from $[N] - L$: namely, the one that cuts all and only the edges (i, j) with $i \in L$ and $j \in V - L$. This cut has total weight

$$S(L) = \sum_{i \in L, j \in V - L} w_{ij}. \quad (3.1)$$

When L consists of more than one vertex, we must have, for any two disjoint contiguous sets L_1, L_2 such that $L = L_1 \cup L_2$,

$$S(L) = S(L_1) + S(L_2) - 2 \sum_{i \in L_1, j \in L_2} w_{ij}. \quad (3.2)$$

This is because the min-cut for L is the sum of the weights of the edges separating L_1 from $[N] - L_1$ and L_2 from $[N] - L_2$, except that the weights of the edges connecting L_1 to L_2 (and L_2 to L_1) must be omitted.

This gives a system of linear equations that can be solved for the weights by induction in the size of L . For example, we have

$$S(AB) = S(A) + S(B) - 2w_{AB}, \quad (3.3)$$

$$S(BC) = S(B) + S(C) - 2w_{BC}, \quad (3.4)$$

and so on. Rearranging, we see that

$$w_{AB} = \frac{S(A) + S(B) - S(AB)}{2}, \quad (3.5)$$

$$w_{BC} = \frac{S(B) + S(C) - S(BC)}{2}, \quad (3.6)$$

and so on; all such weights are nonnegative because of Subadditivity (2.4), which follows from Strong Subadditivity (2.5).

Next consider the size-3 contiguous regions, e.g. ABC :

$$\begin{aligned} S(ABC) &= S(AB) + S(C) - 2w_{AC} - 2w_{BC} \\ &= S(AB) + (S(BC) - S(B) - S(C) + 2w_{BC}) + S(C) - 2w_{AC} - 2w_{BC} \\ &= S(AB) + S(BC) - S(B) - 2w_{AC}, \end{aligned} \quad (3.7)$$

where in the first line we have taken $L_1 = AB, L_2 = C$ in (3.2), and in the second line we have added and subtracted $S(BC)$ using (3.4). The above can be considered as a more general form of (3.2), where L is divided into two *overlapping* regions, $L = L_1 \cup L_2$:

$$S(L) = S(L_1) + S(L_2) - S(L_1 \cap L_2) - 2 \sum_{i \in L_1 - L_2, j \in L_2 - L_1} w_{ij}. \quad (3.8)$$

Formula (3.8) suffices to solve for the remaining edge weights in the graph by expanding the contiguous regions until they reach size $N/2$. For example, we have

$$w_{AC} = \frac{S(AB) + S(BC) - S(B) - S(ABC)}{2}, \quad (3.9)$$

$$w_{AD} = \frac{S(ABC) + S(BCD) - S(BC) - S(ABCD)}{2}, \quad (3.10)$$

and so forth.

All of these expressions for the weights involve only entropies of contiguous regions, which are included in the input data; and all of the weights are nonnegative by Strong Subadditivity (2.5), e.g. $S(ABC) + S(B) \leq S(AB) + S(BC)$. So, since the input data obeys SSA, the N -vertex bulkless graph with weights given by (3.8) solves the DBRP for this data. Furthermore, since each individual weight is fixed by (3.8), the graph is uniquely specified.

Finally, note that we calculated each of the $\binom{N}{2}$ edge weights as a linear combination of at most 4 input entropies. This is trivial to do in $O(N^2)$ time. $\square$

Of course, *every* entropy vector obtained from a quantum state obeys SSA. So one might wonder: how, if at all, does the bulkless graph “remember” that the input vector $v \in \mathbb{R}_{\geq 0}^{\binom{N}{2}}$ came specifically from a *holographic* state?

The answer is interesting and subtle. The bulkless graph indeed solves the DBRP for *arbitrary* SSA-obeying contiguous data. Now, by relations such as (2.7), the bulkless graph *also* solves the DBRP when given as input a full $2^{N-1} - 1$ -dimensional entropy vector —

but only *assuming* the entropies of non-contiguous regions are given by the Bao-Chatwin-Davies formula from section 2.5. In other words, our construction of a bulkless graph doesn't require the contiguous boundary entropies to behave like areas. It doesn't even care about (e.g.) the ordering of N atomic regions around the boundary, except insofar as the ordering affects SSA! On the other hand, as soon as we ask about the entropies of *non-contiguous* boundary regions, our construction will generally return the correct answers only if it was applied to a genuine holographic state. More concretely, as we discussed in section 2.5 above, the Bao-Chatwin-Davies formula will fail when the MMI inequality (2.6) is violated, as occurs in generic non-holographic states.

On the positive side, the bulkless graph has only N vertices and $\binom{N}{2}$ edges, and can be constructed in linear time. On the negative side, it's extremely far from being planar or even “geometric.” However, as the next sections will show, if we use the bulkless graph as a starting point, we can easily construct planar graphs that have the same min-cut structure.

3.2 The chord construction

As figure 8 already suggests, a bulkless graph can be “planarized” — made into a planar graph — by placing the N vertices around a circle, drawing chords for the edges, and then creating a new vertex at each intersection of chords³

Lemma 6 (chord construction). *Given a bulkless graph G with N vertices which solves the DBRP for contiguous input data, there exists a planar graph H with $O(N^4)$ vertices and edges which solves the DBRP for the same data.*

Proof. The proof is essentially pictorial: see figure 9. With the N vertices of the bulkless graph G equally spaced around a circle, we first draw all possible chords between pairs of vertices. We then place a new “bulk” vertex at each intersection of chords. There are $\binom{N}{2}$ chords, and each pair of chords can intersect at most once, so this produces a planar graph H with $O(N^4)$ vertices as well as $O(N^4)$ edges.

All that remains is to choose edge weights for H so that its min-cut structure is identical to that of G . To do this, we simply let the weight of each edge e in H , equal the weight of the edge in G (i.e., the chord) that e came from. For example, all edges along the chord from A to C have weight w_{AC} .

Because all weights in G are nonnegative, so too are all weights in H . Furthermore, given any contiguous boundary region R , any cut in H separating R from $[N] - R$ must intersect every chord of the form (v, w) , for vertices $v \in R$ and $w \notin R$. Now, let C be a cut that itself arises from a chord drawn between R and $[N] - R$. Then C intersects every such chord (v, w) exactly once, and is therefore a min-cut between R and $[N] - R$. Moreover, the weight of this C is just the sum of the weights of all the chords (v, w) that it intersects — the same as the weight of the corresponding min-cut in G . This shows that H has the same min-cut structure as G , and solves the DBRP for the same input data that G did. $\square$

³Chords have recently played a role in quantum gravity, e.g. in the solution of the double-scaled SYK model [26]. They arise here simply for their geometric role — we're drawing straight lines that end on a circle — and it would be surprising if, beyond this obvious fact, our chords have much to do with chords in SYK and JT gravity.




Figure 9. Chord construction for 7 atomic regions. Draw chords for the edges of the bulkless graph, then place a new vertex wherever two chords intersect, to produce a planar graph with $O(N^4)$ vertices and edges. Any given edge e of the bulkless graph gets broken up into many (up to $O(N^2)$) new edges, each of which inherits the same weight $w(e)$ that the parent edge e had.

Note that, just as the bulkless graph for a given number of vertices was universal, so too is the chord construction: besides the number of atomic regions N , all of the solution information is contained in the weights rather than in the graph structure.

The chord construction solves the DBRP, at least in the case of contiguous entropy data that comes from a single 1D boundary. However, it still has the defect that the number of vertices and edges is $\sim N^4$: quadratically greater than the $\sim N^2$ parameters that we started with. So, in the next subsection, we give a different construction of a planar bulk graph, which solves the DBRP using only $\sim N^2$ vertices and edges: the information-theoretic minimum.

3.3 The diamondwork construction

Our final construction of a universal planar bulk graph yields the following:

Theorem 7 (diamondwork graph). *Suppose we are given the $\binom{N}{2}$ entropies for the contiguous boundary regions of a single boundary with N vertices, and suppose these entropies obey SSA. There exists a planar graph which solves the DBRP for this input data with $N^2/2 + O(N)$ vertices and $N^2 + O(N)$ edges. Furthermore, for each N this graph is universal: only the edge weights depend on the input data, not the graph itself or the min-cuts. The edge weights can be computed in $O(N^2)$ time, which is linear in the amount of input data.*

Proof. We will prove theorem 7 in three steps. We first describe how to construct the diamondwork graph and assign weights to its edges in terms of the weights of the bulkless graph (lemma 5) constructed from the same input data. We next show that for every




Figure 10. The diamondwork construction for $N = 7$ atomic regions. The boundary, at the bottom, is a circle, and so the left and right edges of the figure are identified. The k^{th} layer inward from the boundary (k^{th} layer upward in the figure) has one vertex for each contiguous size- k region. Each vertex is connected inwards (up in the figure) to the vertices representing the two contiguous regions of size $k + 1$ it is contained in, and outwards (down in the figure) to the two contiguous regions of size $k - 1$ it contains. For example, the vertex CE , representing the size-3 contiguous region CDE , is connected inwards to BE and CF , and outwards to CD and DE . The graph for N atomic regions has N vertices per layer, and $(N + 1)/2$ layers when N is odd (the case where N is even is slightly more complicated). The weight of an edge is obtained from the weights w_{ij} of the bulkless graph by extending the edge inwards and taking the sum of the weights corresponding to the vertices the extended edge hits. For example, the weight of the edge connecting BC to AC , which extends through AC to CG , is therefore $w_{AC} + w_{CG}$. As argued in the text, this means that min-cuts in the diamondwork graph have the same weights as min-cuts in the bulkless graph. Min-cuts for the boundary regions B , CD , and EFG are shown in the figure.

contiguous region L , there exists a cut in the diamondwork graph with weight $S(L)$. Finally, we show that these cuts are min-cuts, and hence that the diamondwork graph solves the DBRP.

Throughout, we will refer extensively to figures 10 and 11, which depict the diamondwork graph for $N = 7$. We will therefore work directly with the case of N odd, and only sketch the modifications required for the N even case. It will sometimes be convenient to refer to atomic regions by numbers rather than Roman letters, e.g. N denotes the last atomic region.

First we define the diamondwork graph. As shown in figures 10 and 11, the graph contains one vertex for each contiguous boundary region of size at most $N/2$. The vertices are arranged in concentric layers: on the boundary, we have one vertex for each atomic boundary region. Then, one layer inward, there's a vertex for each contiguous boundary




Figure 11. The diamondwork construction for $N = 7$ atomic regions, placed on a circular boundary. The same min-cuts as in figure 10 are shown.

region of size 2, then a vertex for each region of size 3, and so on as we go deeper into the bulk. Each layer is connected to the layers before and after it in a diamondlike pattern (hence the name).

It remains to assign the edge weights. For convenience, we will write weights in terms of the weights w_{ij} in the bulkless graph (lemma 5), which in turn can be written in terms of the contiguous entropies via eq. (3.8). The rule for assigning edge weights is shown pictorially in figure 10, and is this: the weight of any edge e , in the diamondwork graph, is the *sum* of the bulkless weights w_{ij} , over all (i, j) that correspond to the vertices that one reaches when one continues along e toward the center of the diamondwork. For example, the weight of the edge connecting A to AB is $w_{AB} + w_{AC} + \dots + w_{A, (N+1)/2}$. Since the w_{ij} 's are nonnegative, clearly these weights are nonnegative as well.

Equivalently, one can think about the diamondwork graph as a sum of many superimposed triangles T_{ij} , one for each pair of boundary vertices i and j . The three vertices of T_{ij} are the boundary vertex i , the boundary vertex j , and the bulk vertex ij . The weight of any edge e is then just the sum of w_{ij} , over all the triangles T_{ij} to which e belongs.

Next we show that, for every contiguous boundary region L , there exists a cut in the diamondwork graph of weight $S(L)$ — the same as in the bulkless graph.

Recall (eq. (3.1)) that in the bulkless graph, for every contiguous boundary region L , we had

$$S(L) = \sum_{i \in L, j \in [N] - L} w_{ij}. \quad (3.11)$$

Assume for simplicity that $i < j$. Then it can be seen pictorially (figure 10) that the cut C_{ij} that starts to the left of i , heads upwards and rightwards, cutting only left-directed edges until it cuts an edge connected to ij , then heads downwards and rightwards, cutting only right-directed edges, until it ends to the right of j , has precisely this total weight.

We can prove this as follows. We’ve seen that, for every edge e , the weight of e is the sum of w_{ij} , over all triangles $T_{ij} = (i, j, ij)$ that contain e . But this means that the weight of any cut C can be thought of as the sum of w_{ij} over all triangles T_{ij} that C intersects — provided that C intersects each triangle at most once. But the cut C_{ij} defined above does this, because both the cut and the triangles span at most $N/2$ boundary vertices each, so there can be no “wraparound” effects, where a cut intersects the same triangle twice.

Indeed, the cut C_{ij} is *precisely* the unique cut that passes exactly once through every triangle with one vertex in L and one vertex in $[N] - L$. Hence, by eq. (3.11), its weight is exactly $S(L)$.

When N is even, to ensure the above we need to modify the diamondwork construction slightly, as follows. We have $N/2$ “normal” layers, and then one innermost layer, with vertex labels of the form $k, k + N/2$. A triangle $T_{k, k+N/2}$ with a vertex in the innermost layer contributes a weight of $\frac{1}{2}w_{k, k+N/2}$ to each edge e that it contains — so, half of a “normal” triangle’s contribution. Observe that a cut C_{ij} , as defined above, will intersect a triangle of the form $T_{k, k+N/2}$ twice if it intersects it at all, and will pick up half of the $w_{k, k+N/2}$ contribution at each intersection. We thereby preserve the property that the weight of C_{ij} equals the sum $S(L)$ of the appropriate weights from the bulkless graph.

Lastly, we show that these cuts C_{ij} , of weight $S(L)$, are in fact min-cuts of the diamondwork graph. We argued before that any cut separating L from $[N] - L$ must pass at least once through every triangle with one vertex in L and one vertex in $[N] - L$. The cut C_{ij} passes *exactly* once through every such triangle, and passes through no other triangles, so the result is immediate.

As a final observation, we can calculate the relevant sums of w_{ij} ’s by starting with the innermost edges of the diamondwork graph and working our way toward the boundary, maintaining “running totals.” This yields an $O(N^2)$ -time algorithm, which is linear in the amount of input data. Since we observed in lemma 5 that the bulkless weights w_{ij} are themselves also computable in $O(N^2)$ time given the input data, the overall running time is linear. $\square$

To build intuition, we can consider deforming the cut C_{ij} either further into the bulk or closer to the boundary. We can see pictorially that both deformations increase the weight. For example, in figure 10, deforming the cut for CD so it passes above BD results in picking up additional weight $2w_{BD}$, while deforming the cut so it passes below CD adds additional weight $2w_{CD}$. In either case, the cut has been forced to enter, and then leave again, a triangle that does not actually separate vertices in L from vertices in $[N] - L$. The fact that the cut cannot be locally deformed without increasing its weight corresponds to the fact that a geodesic cannot be locally deformed without increasing its length.

3.4 Diamondwork examples

It’s instructive to consider a few examples, to see how varying the edge weights in the otherwise fixed diamondwork graph can encode different bulk geometries.

First consider the situation shown in figure 12, where edges connecting vertices in layer k and layer $k + 1$ have weight $1/k$: in other words, the edge weights encode the harmonic




Figure 12. The $N = 7$ diamondwork graph with harmonic edge weights.

series. As a direct consequence, the min-cut for a contiguous boundary region of size L has a total weight of roughly $2 \ln L$ (the factor of 2 arising because the geodesic needs to go both inward and outward). Notably, this has the same dependence on L as the length of a boundary-anchored geodesic in AdS geometry:

$$|\gamma|_{\text{AdS}} = 2L_{\text{AdS}} \ln L/a, \quad (3.12)$$

with a the UV cutoff, in the limit where $L \ll L_{\text{AdS}}$, as well as the Cardy-Calabrese formula for the entanglement entropy of a small subregion R of size L in a $1 + 1$ -dimensional conformal field theory [27],

$$S(R) = \frac{c}{3} \ln L/a. \quad (3.13)$$

Hence this example correctly reproduces, in our discrete setting, a central feature of AdS/CFT: namely that the larger a region on the boundary, the deeper the RT surface for that region penetrates into the bulk.

Second, consider another highly symmetric situation: a diamondwork graph with all edge weights equal to 1, shown in figure 13. Here the weight of the min-cut is simply equal to the number of edges it cuts, and hence the weight for a boundary region of size L grows *linearly* with L . This is what we would get if, instead of penetrating into the bulk, the geodesics simply hug the boundary. Imagine, for example, that the bulk is almost entirely




Figure 13. The 7-region diamondwork graph with constant edge weights.

filled with a gigantic black hole. In that case, geodesics that start and end at the boundary can never penetrate the event horizon; the best they can do is “go around it the long way”, traversing the thin ring between the event horizon and the boundary. One subtlety here is that, in figure 13, there are *also* min-cuts that penetrate arbitrarily deep into the bulk (indeed there must be, since all the edge weights are equal). To sustain the “black hole” interpretation, we have to say that those cuts are irrelevant; what is relevant is only whether there *exist* cuts that hug the boundary (which there are). A second subtlety is that black holes are usually associated to a nonzero temperature, and to mixed states, while the diamondwork graph as we constructed it always describes a pure state — so if figure 13 represents a black hole, then it’s a degenerate black hole.

In some particularly symmetric cases, we can write down the diamondwork weights directly, without needing to construct an auxiliary bulkless graph and pass through eq. (3.8). For example, consider the case where all k -party contiguous unions of atomic regions have identical entropy S_k , i.e. the entropy of atomic regions is S_1 , of unions of two atomic regions S_2 , etc. Then the diamondwork graph, if it exists, can immediately be seen to have the form shown in figure 14. Nonnegativity of the edge weights is enforced by Subadditivity and Strong Subadditivity: Subadditivity (2.4) implies that

$$S_1 \leq S_2 \leq \cdots \leq S_{k-1} \leq S_k, \quad (3.14)$$




Figure 14. The 7-region diamondwork graph in the symmetric case where all one-party atomic regions have entropy S_1 , all two-party regions have contiguous regions have entropy S_2 , etc.

while Strong Subadditivity (2.5) implies that

$$S_k - S_{k-1} \leq S_{k-1} - S_{k-2} \leq \cdots \leq S_3 - S_2 \leq S_2 - S_1. \quad (3.15)$$

4 The multiple-boundary case

Having solved the case of a single 1D boundary with contiguous entropy data in section 3, it's natural to wonder how far we can generalize our bulk reconstruction methods. For starters, what about the case of *two* 1D boundaries, which are connected via a wormhole? Or for that matter, k 1D boundaries connected via a k -holed sphere (see figure 15)?

A central difficulty, in this case, is simply that the quantum state associated with a single boundary B is no longer necessarily pure — that is, we might have $S(B) > 0$. And we have been unable to generalize the diamondwork construction from section 3 from pure to mixed states.

One way to understand the problem with mixed states is that we no longer have the *complementary recovery* identity $S(R) = S([N] - R)$. Because of this, given N atomic regions on a single boundary, the contiguous entropies now comprise $N^2 - N + 1$ independent real parameters rather than merely $\binom{N}{2}$. Hence, just information-theoretically, any




Figure 15. Two- and three-boundary geometries with wormholes.

diamondwork graph would need to have $\sim N$ layers rather than only $\sim N/2$. But the proof of theorem 7 breaks down once we go beyond $N/2$ layers.

Another way to understand the problem with mixed states is that we can no longer rely only on Strong Subadditivity (2.5) to ensure the existence of a bulk graph. At the least, we also need the Monogamy of Mutual Information (2.6), as shown by the following example of three atomic regions A, B, C in an overall mixed state:

$$S(A) = S(B) = S(C) = 4, \quad (4.1)$$

$$S(AB) = S(BC) = S(AC) = 5, \quad (4.2)$$

$$S(ABC) = 4. \quad (4.3)$$

This example satisfies SSA ($8 < 10$) but violates MMI ($16 > 15$). Since MMI is a cutting-and-pasting inequality (see figure 6), it is satisfied by all graphs, and hence no graph with cut data matching these entropies exists.⁴

4.1 2-boundary counterexamples

Even if the diamondwork construction could be extended to mixed states, we now point out yet a further problem with generalizing to multiple 1D boundaries. Namely, in the multi-boundary case, it seems natural to allow input data about regions that span more than one boundary. As soon as we do, however, we can get bulk graphs that are no longer embeddable onto the appropriate 2D surfaces.

In more detail, suppose we have k circular boundaries $B_1, \dots, B_k$, and the input data consists of $S(R_1 \cup \dots \cup R_k)$ for all k -tuples of contiguous boundary regions $R_1 \subseteq B_1, \dots, R_k \subseteq B_k$. Then already when $k = 2$, we claim that there *cannot* be any solution as simple as our solution to the single-boundary case. We will prove this in two senses:

- (1) Even when there exists a solution, there does not always exist a *planar* solution (when $k = 2$), or more generally, a solution that's embeddable onto a sphere with k holes.

⁴We can purify the system by an additional region D , and think of this data as coming from a *four*-party pure state. Then this example is a manifestation of our statement in section 2.4 that when $N \geq 4$, the holographic entropy cone no longer coincides with the quantum entropy cone, but is a strict subset.




Figure 16. A cutting-and-pasting proof that, if two pairs of straddling geodesics cross each other, then they cannot be minimal surfaces. The four geodesics can be rearranged into four non-minimal surfaces that end on A , A' , B and B' respectively. Hence their combined length must exceed $S(A) + S(B) + S(A') + S(B')$.

- (2) The set of input vectors $v \in \mathbb{R}_{\geq 0}^{\binom{N}{2}^k}$ that do admit planar solutions is not closed under nonnegative linear combinations.

To start with (1), our counterexample will rely on the following lemma.⁵

Lemma 8. *Consider a wormhole with two 1D boundaries. Let A and B be disjoint contiguous regions on the left boundary, and let A' and B' be disjoint contiguous regions on the right boundary. Suppose the only RT surfaces for the regions AB' and BA' both involve ‘straddling’ geodesics, i.e. geodesics that pass between two boundaries. Then those geodesics cannot cross each other.*

Proof. We give a pictorial proof in figure 16. Briefly, we assume by contradiction that the only RT surfaces for AB' and BA' both contain straddling geodesics. We then cut-and-paste those geodesics, rearranging them into non-minimal surfaces for A , B , A' , and B' separately. Taking unions, this gives us non-minimal surfaces for AB' and BA' separately. Since the total area of these surfaces equals the total area of the surfaces that we started with, and since neither new surface contains a straddling geodesic, this gives us our contradiction. $\square$

Using lemma 8, we can now show that, when there are two circular boundaries, it’s no longer true that all input data can be explained by a planar graph if it can be explained by any graph at all.

Theorem 9. *Consider a wormhole with two 1D boundaries. Let A, B, C be disjoint contiguous regions on the left boundary, and let A', B', C' be disjoint contiguous regions on the right boundary. Then there exist entropies for the 9 regions AA' , AB' , AC' , BA' , BB' , BC' , CA' , CB' , CC' that can be explained by some bulk graph, but not by any planar bulk graph (or equivalently, by any bulk graph embeddable onto the wormhole itself).*

⁵A similar construction, forbidding minimal surfaces which cross each other, appeared in section 3.5 of [28]. (We thank an anonymous reviewer for *JHEP* for bringing this construction to our attention.).




Figure 17. The non-planar graph $K_{3,3}$ embedded into a two-boundary geometry, via 9 pairs of straddling geodesics.

Proof. We set

$$S(A) = S(B) = S(C) = S(A') = S(B') = S(C') = 3, \quad (4.4)$$

and

$$\begin{aligned} S(AA') &= S(AB') = S(AC') = S(BA') \\ &= S(BB') = S(BC') = S(CA') = S(CB') = S(CC') = 4. \end{aligned} \quad (4.5)$$

We then, crucially, have

$$S(AA') < S(A) + S(A'), \quad S(AB') < S(A) + S(B'), \quad (4.6)$$

and so on for the other 7 combinations. Now observe that, if there existed an RT surface for AA' that did *not* contain straddling geodesics, then we would necessarily have $S(AA') = S(A) + S(A')$, and likewise for the other combinations. So we conclude that, for each of 9 of the regions AA' , AB' , etc., any RT surface must indeed contain straddling geodesics.

But this means that, whatever our bulk graph, it must give rise to pairs of straddling geodesics in a pattern isomorphic to the complete bipartite graph $K_{3,3}$ (see figure 17). Since $K_{3,3}$ is non-planar, this implies that at least two pairs of straddling geodesics must cross each other in any planar graph model. But this contradicts lemma 8. Hence there is no planar graph model.

On the other hand, the entropies above can easily be completed to a complete list of input entropies for which $K_{3,3}$ itself, with a weight of 1 on each edge, provides a *non*-planar graph model. We simply need to set

$$S(AB) = S(BC) = S(CA) = S(A'B') = S(B'C') = S(C'A') = 6, \quad (4.7)$$

$$S(ABC) = S(A'B'C') = 9, \quad (4.8)$$

$$S(ABA') = S(ABB') = S(ABC') = S(BCA') = S(AA'B') = \dots = 5. \quad (4.9)$$

□

Contrast the above the case of a single boundary, where it surprisingly turned out that every contiguous entropy vector that has any realization at all has a planar realization.

We next prove statement (2).

Theorem 10. *Consider a wormhole with two 1D boundaries. Let there be 3 atomic regions per boundary, labeled A, B, C on the left boundary and A', B', C' on the right boundary. Let an input vector, $v \in \mathbb{R}_{\geq 0}^{63}$, specify $S(R)$ for every nonempty subset $R \subseteq \{A, B, C, A', B', C'\}$. Let V be the subset of input vectors for which there exists a planar graph model. Then V is not closed under convex combinations.*

Proof. We'll define three entropy vectors, $v_1, v_2, v_3 \in \mathbb{R}_{\geq 0}^{63}$. All three will have

$$S(A) = S(B) = S(C) = S(A') = S(B') = S(C') = 1. \quad (4.10)$$

The values of $S(AB)$, $S(BC)$, $S(CA)$, $S(A'B')$, $S(B'C')$, $S(C'A')$, and $S(ABC) = S(A'B'C')$ won't matter for this construction; we can simply choose any values for which planar graph models exist (which is not hard to arrange).

What *does* matter is this:

- v_1 has $S(AA') = S(BB') = S(CC') = 1$,
- v_2 has $S(AB') = S(BC') = S(CA') = 1$,
- v_3 has $S(AC') = S(BA') = S(CB') = 1$,

while in all three cases the “other” straddling entropies are all 2: for example, v_1 has

$$S(AB') = S(BC') = S(CA') = S(AC') = S(BA') = S(CB') = 2. \quad (4.11)$$

We now consider the nonnegative linear combination $v_1 + v_2 + v_3$. This satisfies

$$S(A) = S(B) = S(C) = S(A') = S(B') = S(C') = 3, \quad (4.12)$$

while

$$\begin{aligned} S(AA') &= S(AB') = S(AC') = S(BA') = S(BB') = S(BC') \\ &= S(CA') = S(CB') = S(CC') = 2 + 2 + 1 = 5, \end{aligned} \quad (4.13)$$

which is strictly less than

$$6 = S(A) + S(A') = S(A) + S(B') = \dots. \quad (4.14)$$

Just like in theorem 9, this implies that the 9 regions AA' , AB' , AC' , BA' , BB' , BC' , CA' , CB' , CC' all require pairs of geodesics that straddle the wormhole. Hence these geodesic pairs have the pattern of the non-planar graph $K_{3,3}$ and there must be at least one crossing. So, again applying lemma 8, we find that no graph model for $v_1 + v_2 + v_3$ can be planar. Hence the set V of input data that admits a planar graph model is not closed under convex combination. $\square$

Note that, even with a *single* boundary, if the input data can consist of entropies for non-contiguous regions, then by adapting the above counterexamples — and in particular, by embedding (say) the complete graph K_5 via straddling geodesics — we can create situations where there exists a bulk graph, but there does not exist any *planar* bulk graph. In other words, the key to the above counterexamples is not multiple boundaries *per se*, but rather the non-contiguous input data that multiple boundaries make natural.

4.2 Size upper bound for embeddable graphs

Stepping back, we saw in subsection 4.1 that, when we try to generalize our methods to handle multiple boundaries, a new phenomenon arises: namely, we now encounter geodesics that are not, themselves, the RT surfaces of any boundary region. We called these “straddling geodesics.” These geodesics gave rise to graph models that were no longer embeddable on the relevant 2D surfaces — even by taking convex combinations of input vectors that *do* give rise to embeddable graph models.

Nevertheless, straddling geodesics are still minimal geodesics. As such, they have the crucial property that any two such geodesics intersect in at most one point (see figure 18). By using this property, we can prove an upper bound of $O(N^4)$ on the number of vertices that could ever be needed in a bulk graph, *assuming* that the bulk graph is embeddable onto a 2D surface with one hole for each boundary. This, of course, is dramatically better than the generic 2^{2^N} upper bound on the number of vertices from proposition 2. It’s not as good as the $O(N^2)$ from the diamondwork construction, nor does it lead to a polynomial-time (let alone linear-time) algorithm for *finding* the graph. These we leave as open problems.

Theorem 11 (embeddable graphs need only $O(N^4)$ vertices). *Let $\mathcal{S}$ be a 2D surface with k circular boundaries, which are divided into N atomic regions in total. Given a list of entropies for various unions of those atomic regions (which need not be contiguous), suppose there exists a weighted, undirected graph G solving the DBRP which can be embedded onto $\mathcal{S}$. Then there exists another graph H with at most $O(N^4)$ vertices, which is also embeddable onto $\mathcal{S}$ and which has the same min-cuts as G for all boundary regions.*

Proof. Because G is embeddable onto a 2D surface $\mathcal{S}$, every min-cut can be represented as (possibly a finite union of) curves on $\mathcal{S}$, which we call *geodesics*. It suffices to consider the $\binom{N}{2}$ geodesics that start and end at the points between atomic boundary regions. Just as in the proof of proposition 2, we can reproduce the min-cut structure of G using at most one vertex in each region formed from the intersections of these geodesics. So it suffices to upper-bound the number of such intersection regions.

Recall that we can always choose the geodesics so that every pair crosses each other at most once (see figure 18). Let $M \leq \binom{N}{2}$ be the number of minimally-intersecting geodesics, and imagine drawing them one at a time according to some ordering. Then the t^{th} geodesic can create a new region by intersecting one of the previous $t-1$ geodesics, or by intersecting the boundary at the end of its route. Hence, letting $f(t)$ be the number of regions formed by the first t geodesics, we have the recurrence relation $f(t) \leq f(t-1) + t$, with base case




Figure 18. Ensuring that two geodesics cross at most once. In the left figure, the red and blue curve cross each other multiple times. However, between each crossing, we can pick whichever curve has smaller length between the crossings, and use that for both curves, offsetting them by a slight amount to prevent intersection. This process can be continued until at most one crossing remains. The right figure shows the use of this procedure to eliminate the top two crossings in the left figure.

$f(0) = 1$. The solution is of course

$$f(M) = \frac{M(M+1)}{2} + 1 = O(N^4). \quad (4.15)$$

□

We note that, not by coincidence, the sequence of integers $f(1), f(2), \dots$ from theorem 11 is the so-called “Lazy Caterer’s Sequence”, giving the maximum number of pieces into which a circular cake can be divided using M straight cuts.

From an algorithmic standpoint, the obvious drawback of theorem 11 is that it requires an embeddable bulk graph as its starting point. Nevertheless, theorem 11 is already enough to imply the following:

Corollary 12. *Let $\mathcal{S}$ be a 2D surface with k circular boundaries. Then given as input the entropies $S(R_1), \dots, S(R_L) \geq 0$ for any list of boundary regions $R_1, \dots, R_L$ (which need not be contiguous), the decision problem of whether the $S(R_i)$ ’s admit a graph model that’s embeddable onto $\mathcal{S}$ is in NP (Nondeterministic Polynomial-Time).*

Proof. Note that we can express each R_i as the union of at most $N/2$ atomic boundary regions. So by theorem 11, if a graph model G embeddable onto $\mathcal{S}$ exists, then such a G exists with at most $O(L^4)$ vertices. We are almost done, but to get an NP witness, we also need the technical fact that whenever G exists, we can choose weights for its edges that are expressible with only polynomially many bits each. This follows from the fact that, once we know the set of edges in the min-cut for each R_i , we can express the vector of edge weights as the solution to a linear program, albeit one with $\exp(L)$ constraints. We can then construct a basic feasible solution using Gaussian elimination. □

Note that, whenever we can find disjoint unions of boundary regions that all require straddling geodesics, such that those geodesics have the pattern of $K_{3,3}$ or some other non-planar graph, that constitutes a certificate that there can be *no* planar bulk graph. And likewise for any higher-genus surface $\mathcal{S}$: to construct a no-certificate, it suffices to find straddling geodesics forming a graph that cannot be embedded onto $\mathcal{S}$. However, it is unclear whether such certificates exist whenever there is no embeddable graph, *and* whether they can be found in polynomial time whenever they exist. Note that, if the first statement is true, then the embeddable DBRP is in the class $\text{NP} \cap \text{coNP}$; if both are true, then the problem is in P .

5 Discussion and open problems

Having presented our main results, we conclude with some possible generalizations and open research directions.

Higher dimensions. The immediate obstacle we face when the boundary has more than one spatial dimension is that there is now an *exponential*, rather than polynomial, number of contiguous regions. It is an interesting question whether there is a different, polynomially-sized set of entropic data which suffices, in some cases, to reproduce the rest of the entropy — that is, is there a higher-dimensional analogue of eq. (2.7)? One obvious candidate on a 2-dimensional boundary would be the set of discs of a fixed radius. However, it is unclear what the appropriate generalizations of our bulkless, chord, or diamondwork constructions would be. Perhaps we no longer want a graph at all, but a *simplicial complex*.

Beyond RT: dynamics and corrections. Note that (2.7) requires the RT formula to hold *exactly*; if we create, for example, an EPR pair shared between A and C , it is clear that the left hand side will decrease while the right hand side will not change. (If we create *many* EPR pairs, the ER=EPR conjecture [29] suggests that we will form a wormhole, and thus be back in the purview of section 4.) All of our graph constructions should therefore be understood to be discretizations of fully classical geometries. Furthermore, we have tacitly assumed that all of the extremal surfaces lie on a single timeslice of the bulk. In general, this need not be the case: there exist dynamical spacetimes, such as collapsing shockwave geometries, where extremal surfaces are not spacelike. In cases like this, we should use not the RT formula but the HRT formula [28, 30], which correctly prescribes the appropriate way to extremize over such surfaces. We did not consider such dynamical spacetimes further in the body of the paper, where we were really trying to reconstruct a discrete bulk geometry at one point in time rather than a full spacetime. There doesn't seem to be any fundamental obstacle to trying to discretize a spacetime rather than a graph, but it would have to be a discretization of at least a three-dimensional manifold, with all the difficulties discussed before. For some existing work on entropy relations and discretizations for dynamical spacetimes, see [15, 31].

Beyond min-cuts. We commented at the outset that in some sense the DBRP is the *inverse* of the classical min-cut problem: instead of being given a fixed graph and asked to

compute the min-cut separating two sets of vertices, we are given cut-values and asked to construct a graph. It would be interesting to understand the relation of the DBRP to other min-cut (and max-flow) problems. For example, in the multiflow literature, a flow is said to *lock* a region if it saturates the bound given by the max-flow/min-cut theorem, and various “locking theorems” guarantee the existence of flows that lock all contiguous boundary regions simultaneously [32, 33]. Is there any relation to our results, which also involve a special role for contiguous boundary regions? Also, once we’re considering max-flows in addition to min-cuts, do the so-called *quantum max-flows* (e.g., [34]) play any interesting role in AdS/CFT, and if so can we construct bulk graphs that account for them?

More about multiple boundaries. We showed in section 4 that, when straddling geodesics exist, in general there is no planar graph that reproduces the boundary entropies. Nevertheless, theorem 11 tells us that there exists a polynomial set of input data — the lengths of *all* N^2 geodesics, both RT surfaces and those passing between boundaries — which suffices to reproduce the entropy vector. One obstruction to passing from this observation to a graph is that it is unclear how to derive the lengths of geodesics that are not themselves RT surfaces from the entropy vector. Can this be done? Along similar lines, in section 4 we worked with the $O(N^{2k})$ entropies of unions of contiguous regions. Yet theorem 11 again shows that embeddable graphs, which include those with the topology of a k -punctured sphere, need at most $O(N^4)$ vertices, independent of the value of k . Does this mean there is a smaller set of input data — for example, the entropies of all *pairs* of contiguous boundary regions — that already suffices to determine the entropies of *all* boundary regions? If so, is there some way to use that input data to construct versions of bulkless or diamondwork graphs that work for multi-boundary geometries, or even just a single boundary in a mixed state?

From weights to geometry. As shown in section 3, our various graph constructions (bulkless, chord, diamondwork) for a single boundary encode all of the geometric information in the edge weights and none in the graph structure: that is what it means for them to be “universal.” To think of the graph as an actual discrete geometry, we might instead want to move the geometric information from the weights to the graph structure. Forthcoming work by one of us and collaborators [35] presents an approximate numerical algorithm to accomplish this task. Furthermore, once we have a graph it’s natural to wonder whether we can use it, not just to compute entropies, but to define a discretized version of a quantum state, by putting a tensor network on top of the graph. Of our various graph constructions, the diamondwork construction, which has the information-theoretically minimal number of required vertices, seems most suitable to host a tensor network, and it would be interesting to investigate this further.

Graphs beyond holography. Because our graph constructions take only the contiguous entropies as input, they can be applied even to entropy vectors *outside* the holographic entropy cone. In that case, however, our constructions would produce graphs that generate “wrong” predictions for the entropies of some non-contiguous regions. Can we understand the failure of those predictions in any quantitative detail? What about for entropy vectors

that are outside the holographic entropy cone, but still in the stabilizer entropy cone (see, e.g., [36])? More generally, what insight can we glean from our graph constructions about the holographic entropy cone itself?

Recovering the quantum state. In general, an entropy vector represents an equivalence class of quantum states living on N -partite Hilbert spaces, all of which share the same von Neumann entropies of their reduced density matrices. Given an entropy vector, can we efficiently prepare a state $|\psi\rangle$ with the appropriate entropies? In other words, can we go in the opposite direction from this paper: from entropy vector to boundary, rather than from entropy vector to bulk? Can we at least do this for matching vectors, in which all entropies can be obtained from the contiguous ones? There are existing constructions in the literature to provide CFT states [6] and stabilizer tensor networks [10] which automatically satisfy the RT formula for a given set of boundary subregions, but both require an infinitely large Hilbert space. Can we do better than this? We note that it's not hard to construct a state $|\psi\rangle$ that reproduces the entropies of the *contiguous* regions, by starting with the bulkless graph, with edge weights $\{w_{ij}\}$, and then inserting w_{ij} EPR pairs (or fractional EPR pairs) between the atomic regions i and j , for all i, j . However, this state will correspond to a geometry with numerous wormholes, and will yield wildly wrong predictions for the entropies of *non*-contiguous regions.

Unordered regions. In this paper, we studied how to produce a bulk graph given entropies of contiguous 1D boundary regions — but always under the assumption that we knew the *order* of atomic regions around the boundary, as well as which entropies correspond to which boundary regions. What if we no longer know one or both of those things, and are just handed a “bag of entropies” to make sense of? Do we then get an NP-hard reconstruction problem? Or possibly something equivalent to the so-called *turnpike problem* in computer science [37]?

Acknowledgments

Scott Aaronson is supported by a Vannevar Bush Fellowship from the U.S. Department of Defense, the Berkeley NSF-QLCI CIQC Center, a Simons Investigator Award, and the Simons “It from Qubit” collaboration.

Jason Pollack is supported by the Simons Foundation through “It from Qubit: Simons collaboration on Quantum Fields, Gravity, and Information”.

The authors thank Ning Bao, Charles Cao, Netta Engelhardt, Matt Headrick, Sergio Hernandez-Cuenca, Alexander Jahn, Greg Kuperberg, Sarah Racz, Vincent Steffan, Lenny Susskind, and Benson Way for helpful discussions and comments.

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. SCOAP³ supports the goals of the International Year of Basic Sciences for Sustainable Development.

References

- [1] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Phys. Rev. Lett.* **96** (2006) 181602 [[hep-th/0603001](#)] [[INSPIRE](#)].
- [2] N. Engelhardt and A.C. Wall, *Extremal Surface Barriers*, *JHEP* **03** (2014) 068 [[arXiv:1312.3699](#)] [[INSPIRE](#)].
- [3] 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](#)].
- [4] N. Bao, C. Cao, S. Fischetti, J. Pollack and Y. Zhong, *More of the Bulk from Extremal Area Variations*, *Class. Quant. Grav.* **38** (2021) 047001 [[arXiv:2009.07850](#)] [[INSPIRE](#)].
- [5] P. Hayden, M. Headrick and A. Maloney, *Holographic Mutual Information is Monogamous*, *Phys. Rev. D* **87** (2013) 046003 [[arXiv:1107.2940](#)] [[INSPIRE](#)].
- [6] 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](#)].
- [7] A. Almheiri, X. Dong and D. Harlow, *Bulk Locality and Quantum Error Correction in AdS/CFT*, *JHEP* **04** (2015) 163 [[arXiv:1411.7041](#)] [[INSPIRE](#)].
- [8] D. Harlow, *The Ryu-Takayanagi Formula from Quantum Error Correction*, *Commun. Math. Phys.* **354** (2017) 865 [[arXiv:1607.03901](#)] [[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] N. Bao and A. Chatwin-Davies, *The Complexity of Identifying Ryu-Takayanagi Surfaces in AdS_3/CFT_2* , *JHEP* **11** (2016) 034 [[arXiv:1609.01727](#)] [[INSPIRE](#)].
- [12] L.R. Ford and D.R. Fulkerson, *Maximal flow through a network*, *Can. J. Math.* **8** (1956) 399.
- [13] T.H. Cormen, C.E. Leiserson, R.L. Rivest and C. Stein, *Introduction to algorithms*, MIT Press (2022).
- [14] L. Chen, R. Kyng, Y.P. Liu, R. Peng, M.P. Gutenberg and S. Sachdeva, *Maximum Flow and Minimum-Cost Flow in Almost-Linear Time*, [arXiv:2203.00671](#).
- [15] M. Rota and S.J. Weinberg, *New constraints for holographic entropy from maximin: A no-go theorem*, *Phys. Rev. D* **97** (2018) 086013 [[arXiv:1712.10004](#)] [[INSPIRE](#)].
- [16] B. Czech and X. Dong, *Holographic Entropy Cone with Time Dependence in Two Dimensions*, *JHEP* **10** (2019) 177 [[arXiv:1905.03787](#)] [[INSPIRE](#)].
- [17] T. He, M. Headrick and V.E. Hubeny, *Holographic Entropy Relations Repackaged*, *JHEP* **10** (2019) 118 [[arXiv:1905.06985](#)] [[INSPIRE](#)].
- [18] N. Bao, N. Cheng, S. Hernández-Cuenca and V.P. Su, *A Gap Between the Hypergraph and Stabilizer Entropy Cones*, [arXiv:2006.16292](#) [[INSPIRE](#)].
- [19] M. Fadel and S. Hernández-Cuenca, *Symmetrized holographic entropy cone*, *Phys. Rev. D* **105** (2022) 086008 [[arXiv:2112.03862](#)] [[INSPIRE](#)].

- [20] A. Bouland, B. Fefferman and U. Vazirani, *Computational pseudorandomness, the wormhole growth paradox, and constraints on the AdS/CFT duality*, [arXiv:1910.14646](#) [[INSPIRE](#)].
- [21] A.R. Brown, H. Gharibyan, G. Penington and L. Susskind, *The Python's Lunch: geometric obstructions to decoding Hawking radiation*, *JHEP* **08** (2020) 121 [[arXiv:1912.00228](#)] [[INSPIRE](#)].
- [22] N. Engelhardt, G. Penington and A. Shahbazi-Moghaddam, *Finding pythons in unexpected places*, *Class. Quant. Grav.* **39** (2022) 094002 [[arXiv:2105.09316](#)] [[INSPIRE](#)].
- [23] C. Kuratowski, *Sur le probleme des courbes gauches en topologie*, *Fund. Math.* **15** (1930) 271.
- [24] M. Walter, *Multipartite Quantum States and their Marginals*, Ph.D. Thesis, ETH Zurich, Zurich, Switzerland (2014) [[DOI:10.3929/ethz-a-010250985](#)] [[arXiv:1410.6820](#)] [[INSPIRE](#)].
- [25] M. Walter and F. Witteveen, *Hypergraph min-cuts from quantum entropies*, *J. Math. Phys.* **62** (2021) 092203 [[arXiv:2002.12397](#)] [[INSPIRE](#)].
- [26] M. Berkooz, M. Isachenkov, V. Narovlansky and G. Torrents, *Towards a full solution of the large N double-scaled SYK model*, *JHEP* **03** (2019) 079 [[arXiv:1811.02584](#)] [[INSPIRE](#)].
- [27] P. Calabrese and J.L. Cardy, *Entanglement entropy and quantum field theory*, *J. Stat. Mech.* **0406** (2004) P06002 [[hep-th/0405152](#)] [[INSPIRE](#)].
- [28] 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](#)].
- [29] J. Maldacena and L. Susskind, *Cool horizons for entangled black holes*, *Fortsch. Phys.* **61** (2013) 781 [[arXiv:1306.0533](#)] [[INSPIRE](#)].
- [30] V.E. Hubeny, M. Rangamani and T. Takayanagi, *A Covariant holographic entanglement entropy proposal*, *JHEP* **07** (2007) 062 [[arXiv:0705.0016](#)] [[INSPIRE](#)].
- [31] A. May, *Tensor networks for dynamic spacetimes*, *JHEP* **06** (2017) 118 [[arXiv:1611.06220](#)] [[INSPIRE](#)].
- [32] A. Frank, A.V. Karzanov and A. Sebo, *On Integer Multiflow Maximization*, *SIAM J. Discrete Math.* **10** (1997) 158.
- [33] M. Headrick, J. Held and J. Herman, *Crossing Versus Locking: Bit Threads and Continuum Multiflows*, *Commun. Math. Phys.* **396** (2022) 265 [[arXiv:2008.03197](#)] [[INSPIRE](#)].
- [34] S.X. Cui, M.H. Freedman, O. Sattath, R. Stong and G. Minton, *Quantum Max-flow/Min-cut*, *J. Math. Phys.* **57** (2016) 062206 [[arXiv:1508.04644](#)] [[INSPIRE](#)].
- [35] C. Cao, J. Pollack and B. Way, forthcoming work (2022).
- [36] C. Keeler, W. Munizzi and J. Pollack, *Entropic lens on stabilizer states*, *Phys. Rev. A* **106** (2022) 062418 [[arXiv:2204.07593](#)] [[INSPIRE](#)].
- [37] T. Dakic, *On the turnpike problem*, Simon Fraser University (2000).