

Topological link models of multipartite entanglement

Ning Bao¹, Newton Cheng², Sergio Hernández-Cuenca³, and Vincent Paul Su²

¹Computational Science Initiative, Brookhaven National Lab, Upton, NY, 11973, USA

²Center for Theoretical Physics, Department of Physics, University of California, Berkeley, CA 94720, USA

³Department of Physics, University of California, Santa Barbara, CA 93106, USA

We introduce a novel model of multipartite entanglement based on topological links, generalizing the graph/hypergraph entropy cone program. We demonstrate that there exist link representations of entropy vectors which provably cannot be represented by graphs or hypergraphs. Furthermore, we show that the contraction map proof method generalizes to the topological setting, though now requiring oracular solutions to well-known but difficult problems in knot theory.

1 Introduction

Originally motivated by relations obeyed by the entropy function in holography [1], [2, 3, 4] studied a graph model of entanglement, wherein the von Neumann entropies of quantum subsystems are calculated by minimum cuts separating those subsystems from their complements in a corresponding weighted graph. The combinatorial properties of minimum cuts on graphs allowed for the development of a method to prove entropy inequalities obeyed by holographic states. While powerful, this technology was limited to graph models intended to describe the subset of holographic states, and therefore did not capture the entanglement structure of more general quantum states. In follow-up work, the graph model was generalized to a hypergraph model [5], with the aim of reproducing the entropies of a wider class of states while retaining the ability to prove entropy inequalities via a generalization of the proof technology. For four parties, this was indeed the case: the hypergraph cone coincided with that of stabilizer states, which is larger than the holographic one of graphs, but still not large enough to capture general quantum states.¹

The hypergraph model also provided interesting new results intimately related to the properties of ranks of linear subspaces and the entropies of stabilizer states in quantum mechanics [5]. Further explorations of their connection to the latter revealed that the hypergraph model was able to encode the entropies of a subset of stabilizer states [8, 9], and that this containment was proper [10]. The latter result relied on the discovery of a particular stabilizer state whose entropy vector, referred to as ray 15 in [10], violates an inequality proven to be valid for all hypergraphs in the hypergraph model.

Ning Bao: ningbao75@gmail.com

Newton Cheng: newtoncheng@berkeley.edu

Sergio Hernández-Cuenca: sergiohc@ucsb.edu

Vincent Paul Su: vipasu@berkeley.edu

¹The full four-party quantum entropy cone is unknown [6], though it is known to properly contain the stabilizer entropy cone as defined in [7].

By contrast, another avenue of study has been the study of topology as an efficient encoding of certain patterns of entanglement. Some work in this direction include [11, 12, 13, 14, 15, 16]. Because all graph and hypergraph models of entanglement can be captured by topological links (as will be discussed in detail later in this work), using topological links to generalize encodings of entanglement patterns in a systematic way is an intriguing avenue of exploration to describe the entanglement of more general quantum states.

In this paper, we introduce such a generalization of the graph and hypergraph models of quantum entanglement based on topological links. Similar to the (hyper)graph case, entropies will be given by the weights of minimal cuts on links, and used to define the *entropy cone of link models* as the set of all entropy vectors realizable by such models. In Section 2, we review the graph and hypergraph models and introduce the contraction map method that was applied to prove entropy inequalities on such models. In Section 3, we then formally introduce link models as a generalization of hypergraph models. We demonstrate that such models strictly contain the hypergraph entropy cone by exhibiting a link whose entropies capture those of ray 15. We then prove a generalized contraction map condition on such models that, subject to access to an oracle whose abilities we carefully define, provides a sufficient condition for the validity of a given entropy inequality on a link model. We also motivate that link models may usefully capture the entropies of quantum states by showing that the entropy cones of link models and quantum states are identical for $n \leq 3$ parties. A more detailed discussion of how link models can be applied to further characterize entanglement structures of quantum states is provided in Section 4.

2 Preliminaries

2.1 Graph models

We first review the definition of the graph model and the intuition behind it, as originally introduced and studied in [2]. For any positive integer k , let $[k] = \{1, \dots, k\}$. A graph model consists of an undirected graph $G = (V, E)$ with nonnegative edge weights specified by a map $w : E \rightarrow \mathbb{R}_{\geq 0}$. A subset of vertices $\partial V \subseteq V$ are called *external vertices* and the rest are referred to as *internal vertices*. For ∂V of cardinality $n + 1$, a bijective coloring $\partial V \rightarrow [n + 1]$ is used to identify internal vertices with partitions of a quantum state: G represents the entropies of a mixed n -party state on the vertices $[n]$, with vertex $n + 1$ representing the purification of the state, and all internal vertices encoding the entanglement structure. In this setting, there are $2^n - 1$ subsystems of the quantum state one can consider, corresponding to every non-empty subset $I \subseteq [n]$. The entropy of any one of these is defined as the weight of the minimum cut or *min-cut* separating vertices I from $\partial V \setminus I$.

Given an arbitrary graph model, one may intuitively expect that the min-cut prescription should place restrictions on the entropies of states that can be faithfully represented by such models. These restrictions turn out to take the form of linear entropy inequalities, which can be canonically written as

$$\sum_{l=1}^L \alpha_l S(I_l) \geq \sum_{r=1}^R \beta_r S(J_r), \quad (2.1)$$

where $\alpha_l, \beta_r > 0$, and I_l, J_r are subsets of external vertices. In other words, we arrange terms into a left-hand side (LHS) and a right-hand side (RHS) such that they all appear with positive coefficients. The entropies of any state which can be represented by a graph model will obey inequalities of the general form of (2.1).

These inequalities on graph models can be proven via a technique known as proof-by-contraction. The method functions by essentially formalizing the intuition that the weights of any clever construction of cuts corresponding to a given choice of external vertices are always lower bounded by the weights of the corresponding min-cut.² To give the precise statement of the proof-by-contraction method, let us first introduce some basic notation. For each $i \in [n + 1]$, let $x^{(i)}$ and $y^{(i)}$ respectively be bitstrings of lengths L and R , with entries given by $x_l^{(i)} = \delta(i \in I_l)$ and $y_r^{(i)} = \delta(i \in J_r)$, where δ is a Boolean indicator function yielding 1 or 0 depending on whether its argument is true or false, respectively. These are often referred to as *occurrence bitstrings*, as they simply encode the occurrence of the i^{th} party in I_l or J_r . Given a vector $\gamma \in \mathbb{R}^k$, we also introduce a *weighted Hamming norm* on \mathbb{R}^k as $\|x\|_\gamma = \sum_{j=1}^k \gamma_j |x_j|$. We can now state the proof-by-contraction method, a sufficient condition for the validity of a given entropy inequality [2]:

Theorem 1. *Let $f : \{0, 1\}^L \rightarrow \{0, 1\}^R$ be a $\|\cdot\|_\alpha$ - $\|\cdot\|_\beta$ contraction, i.e.*

$$\|x - x'\|_\alpha \geq \|f(x) - f(x')\|_\beta \quad \forall x, x' \in \{0, 1\}^L. \quad (2.2)$$

If $f(x^{(i)}) = y^{(i)}$ for all $i = 1, \dots, n + 1$, then (2.1) is a valid entropy inequality on graphs.

The final requirement on occurrence bitstrings is a simple consistency condition that only the appropriate external vertices from LHS min-cuts are included by f in forming cuts for each RHS set J_r . If f is a function satisfying the contraction condition in Theorem 1, we call f a *contraction map*. For the detailed proof, we refer the interested reader to [2, 4], but we briefly review the elements here. At the core of the proof is constructing a partitioning of the underlying graph into 2^L subsets of vertices as determined by inclusion/exclusion with respect to each of the L min-cuts of LHS terms in (2.1). Each of these 2^L subsets is conveniently labeled by a bitstring $x \in \{0, 1\}^L$. Given a bitstring x , one can reconstruct the associated vertex subset via $W(x) = \bigcap_{l=1}^L W_l^{x_l}$, where $W_l^1 = W_l$ is the min-cut for the l^{th} region, and $W_l^0 = W_l^c$ is its complement. This may also be inverted to recover the cuts themselves via $W_l = \bigcup_{x_l=1} W(x)$, where the union runs over all $x \in \{0, 1\}^L$ subject to $x_l = 1$. While the entropy will depend on the actual weights of the edges in a graph, the construction of these subsets $W(x)$ allows one to keep track of how many times we count edge weights coming out of that subset when computing entropies.

The contraction map f takes bitstrings of length L to bitstrings of length R , and in this way utilizes the subsets $W(x)$ to build (not necessarily minimal) cuts U_r for each of the RHS subsystems. More explicitly, these cuts are specified in terms of f by $U_r = \bigcup_{f(x)_r=1} W(x)$ (cf. the reconstruction of the LHS cuts W_l above). By minimality of cuts, the sum of edge weights involved in every U_r will upper bound the true entropy of terms on the RHS. Thus, f can be thought of as giving a sort of naive overestimate of the entropies of terms on the RHS, as long as the cuts U_r are appropriately faithful (e.g., vertices in a cut for I are not all excluded from vertices in a cut for $I \cup J$).

The contraction condition on f , while looking cryptic and abstract since it is defined on bitstrings, enforces that edges coming out of the $W(x)$ are counted more times on the LHS of the inequality than on the RHS. Because it holds for every pair of bitstrings, which correspond to different sets $W(x)$, it says that edges between different $W(x)$ were counted no more times on the RHS than on the LHS, hence proving the inequality (2.1) valid regardless of the edge weights in the underlying graph. Determining the existence of a

²See [17] for a description of the contraction-map proof method in the geometric setting that originally inspired it, based on the holographic Ryu-Takayanagi formula [1].

contraction map for a given entropy inequality can be accomplished via greedy algorithms which are guaranteed to find one if it exists [2]. Unfortunately though, this rapidly becomes computationally intractable, since one is effectively searching in a space of all possible $2^{R \times (2^L - (n+1))}$ maps f compatible with the last line of Theorem 1 as specified by their image bitstrings.³

2.2 Generalizing to hypergraphs

In [5], this was generalized to hypergraph models, where, again, entropies were declared to be calculated by weights of min-cuts. The sets $W(x)$ were again useful, but because hyperedges can connect more than two vertices at a time, one had to consider not just pairwise interactions between the $W(x)$, but also general k -edges, i.e., hyperedges connecting k vertices. Demanding that k -edges on the LHS were counted no fewer times than on the RHS, essentially imposed additional criteria on the contraction map f , which neatly nested into a sum over the rank k of the involved hyperedges. This is captured in a generalization of the contraction condition:

Theorem 2. *Let $f : \{0, 1\}^L \rightarrow \{0, 1\}^R$ be an i_α^k - i_β^k contraction:*

$$i_\alpha^k(x^1, \dots, x^k) \geq i_\beta^k(f(x^1), \dots, f(x^k)) \quad \forall x^1, \dots, x^k \in \{0, 1\}^L. \quad (2.3)$$

If $f(x^{(i)}) = y^{(i)}$ for all $i \in [n + 1]$, then (2.1) is a valid inequality on rank- k hypergraphs.

Here, $i_\alpha^k(x^1, \dots, x^k) = \sum_l \alpha_l i^k(x_l^1, \dots, x_l^k)$ and $i^k : \{0, 1\}^k \rightarrow \{0, 1\}$ is an indicator function that returns 0 if the input bitstring is all 0s or all 1s, and 1 otherwise. Because a valid contraction map f for hypergraphs must obey additional constraints, an inequality that is true for graphs may no longer hold for hypergraphs. Thus, the hypergraph entropy cone has fewer constraints and is larger. That this containment is strict was shown in [5].

3 Link models

3.1 Introduction

Inspired by the previous models, we now introduce the link model. In this paper, a *link* consists of standard 1-dimensional *knots* in 3-dimensional Euclidean space. For our purposes, knots will often be trivial (the *unknot*) and just referred to as *loops*. Letting \mathcal{L} denote the finite set of loops, we distinguish a subset $\partial\mathcal{L} \subseteq \mathcal{L}$ as *external* and call the others *internal*. All loops are assigned some nonnegative weight. Notice that links no longer carry distinct vertex/edge-like components – loops will implement both functionalities now. In analogy with previous discussions, a bijection $\partial\mathcal{L} \rightarrow [n + 1]$ once again identifies the external loops as parties of a quantum state.

To define the notion of loop cuts on a link model in analogy with edge cuts on graph models, we will need the following basic idea:

Definition 1. *An n -loop link L embedded in \mathbb{S}^3 is said to be topologically linked if for any choice of embedded disks $\{D_i \subset \mathbb{S}^3\}_{i=1}^n$ such that $L = \bigcup_{i=1}^n \partial D_i$, the set $\bigcup_{i=1}^n D_i$ is topologically connected. Otherwise L is said to be topologically unlinked.*

³See [4] for a more refined study of the space of possible contraction maps.

Given a link with loop set \mathcal{L} , any subset of loops $\tilde{\mathcal{L}} \subseteq \mathcal{L}$ defines a *sublink*. Similar to the notion of connected components in a topological space, one easily sees that a general link L may always be decomposed into *connected sublinks*, namely, maximal sublinks which are topologically linked. Hence we can equivalently understand a link L as being topologically unlinked if and only if it consists of more than one connected sublink. Clearly, topological linking is a property which can be lost when some subset of loops is removed from a link, in the sense that the remaining loops can break into multiple connected sublinks. Following this intuition, we will say that two loops of a link L are topologically unlinked in L if they belong to distinct connected sublinks, and similarly that two subsets of loops are topologically unlinked in L if every pair of respective loops belong to distinct connected sublinks.⁴

We use the above to define a *loop cut* C for a subsystem $I \subseteq [n]$ as any set of internal loops $C \subseteq \mathcal{L} \setminus \partial\mathcal{L}$ which, if removed, give a link where the loops in the subsystem I are topologically unlinked from all other external loops in $\partial\mathcal{L}$. The weight of C is given by the sum of its loop weights, and denoted by $|C|$. The entropy of I is defined as the minimum weight of all possible loop cuts C for I . Any such *loop min-cut* for I will be denoted by C_I .⁵ In the spirit of previous work, we posit that the entropy of a subsystem I is given by the weight of a loop min-cut C_I . Note that both the graph and hypergraph models are easily reproducible in the link model by converting vertices to infinite-weight loops and k -edges to finite weight loops, all Hopf-linked to their vertex loops, as illustrated in Figure 1.

3.2 Capturing the entropies of ray 15

To motivate the construction of the link model, we share an explicit quantum state whose subsystem entropies can be represented by the link model yet provably fail to be realizable using hypergraph models [5, 10]. The pertinent entropy vector, henceforth referred to as ray 15 following [10], comes from a particular 5-party mixed graph state,⁶ where the vertices denote spin systems and edges represent Ising-type interactions [18]. Explicitly, ray 15 reads

$$\vec{S}_{15} = (11111; 1222222221; 2222222222; 22122; 1), \quad (3.1)$$

where entropies are listed in lexicographical order, with subsystems of increasing number of parties separated by a semicolon. Notably, all graph states are stabilizer states, and hence, ray 15 must lie within the 5-party stabilizer entropy cone.

In attempting to determine whether the hypergraph and stabilizer entropy cones coincide, the authors found that the answer is negative [10], indicating the strict containment of the hypergraph entropy cone within the stabilizer entropy cone [8]. Indeed, ray 15

⁴The converses are clearly as follows: two loops of a link L are topologically linked in L if they belong to the same connected sublink, and similarly two subsets of loops are topologically linked in L if all their loops belong to the same connected sublink. However, let us emphasize an important subtlety: two loops of a link L which are topologically linked in L need not be so if taken to define a link on their own. For instance, if L is a Brunnian link, any two loops in it are topologically linked in L , but if taken to define a link on their own (cf. deleting the other loop), they would be topologically unlinked (indeed, they would define a 2-component unlink). We will come back to this in Section 3.3.

⁵When there exist multiple min-cuts with the same cut weight, one simply makes a choice and sticks to it.

⁶Note that the graph state described here is in the context of [18], and not a state whose entropy vector can be realized by graph representation.




Figure 1: Link construction from graph and hypergraph models. One can reproduce all min-cut weights appropriately by making the following substitutions. Vertices do not contribute to the cut weight, so they should be replaced with loops of infinite weight. Loops of finite weight can be used to replace both edges and hyperedges where the vertices are Hopf-linked to the edge loop. Note that link models inherently no longer distinguish between vertices and edges.




Figure 2: The topological link whose entropies reproduce ray 15 from [10]. A , B , C , D , and E are the parties in question, with F taken to be the purifier. The min-cut of any single external loop has weight 1, as do the pairs AB , ED , and CF . All other pairs of external loops have min-cut weight 2, and so do all triples. The min-cut weight of any other subsystem can be obtained from that of its complement with respect to all external loops.

violates a linear entropy inequality that all hypergraph models were shown to obey in [10]:

$$S_{AB}+S_{DE}+S_{ACD}+2S_{ACE}+S_{BCD}+S_{ABDE} \geq S_{AC}+S_{AE}+S_{BD}+2S_{ABCD}+S_{ACDE}. \quad (3.2)$$

In Figure 2, we give an explicit link model that reproduces ray 15, thereby demonstrating that there exist link models which reproduce entropy vectors inaccessible by hypergraphs alone. Moreover, because all hypergraph models can be reproduced in the link model, the entropy cone of links properly contains the hypergraph entropy cone.

This construction motivates us to investigate the degree to which link models can represent multipartite entanglement of generic quantum states. We will consider whether the contraction-map technology can be generalized to link models, allowing us to extend the entropy cone of hypergraphs to that of link models. This will pave the way for exploring whether the link model entropy cone can possibly contain the stabilizer entropy cone of [7] or even the full quantum entropy cone [6]. One can consider a program of using contraction maps to directly prove candidate entropy inequalities on link models, and exploring such inequalities as candidate entropy inequalities for the currently unknown higher-party stabilizer and quantum entropy cones (e.g. by testing whether they are violated by the entropies of known stabilizer or quantum states). In this way, the contraction map methods and associated graph/link models serve as an immediate application and useful set of tools in the study of quantum state entropies and their constraints. We further discuss the application of link models to classifying the structure of entanglement of quantum states in Section 4.

3.3 The link model toolkit

Studying general link models will require the introduction of various definitions and constructs, to which we now turn. Given some link, we say that a sublink $\tilde{\mathcal{L}}$ of it is an *irreducible sublink* if it is intrinsically topologically linked, i.e. if it is topologically linked once all loops in $\mathcal{L} \setminus \tilde{\mathcal{L}}$ are deleted.⁷ A simple class of irreducible sublinks are the connected sublinks into which any link may be decomposed.

Any loop cut C for I naturally defines two additional subsets of loops. Let $W_C \subseteq \mathcal{L} \setminus C$ be the minimal-cardinality subset of loops with $W_C \cap \partial\mathcal{L} = I$ and such that, once the loops in C are removed, form a collection of connected sublinks. In practice, W_C can be found by collecting all loops in the link defined by $\mathcal{L} \setminus C$ that are topologically linked to the pertinent external loops I . The resulting set W_C is clearly unique and referred to as the *cut interior* of C . Similarly, $\overline{W}_C = \mathcal{L} \setminus (W_C \cup C)$ is also unique and referred to as the *cut exterior* of C . The cut exterior contains all external loops $\partial\mathcal{L} \setminus I$ which are absent from W_C , and consists of connected sublinks as well. When talking about a loop min-cut C_I , we will denote its associated cut interior as W_I for simplicity. As an example, consider looking at subsystem $I = AC$ in Figure 2: a valid loop min-cut C_{AC} for it is the weight-2 loop which, when removed, defines the cut interior $W_{AC} = \{A, C\}$ alone – the corresponding cut exterior \overline{W}_{AC} consists of the three weight-1 loops and the complementary external loops B, D, E and F .

To refine our analysis of loop cuts, we introduce a notion of *prunability*. For a given subsystem I , the loop min-cut C_I by definition contains loops which topologically unlink W_I from \overline{W}_I . Consider a sublink $B_I \subseteq \mathcal{L}$ containing k loops that is irreducible and has nontrivial intersection with all of W_I, \overline{W}_I , and C_I . Let us refer to such a sublink as a *bridge* across the loop min-cut C_I . By definition, bridges clearly have cardinality $k \geq 3$.

⁷Note that a sublink of some link L may not be irreducible even if its constituent loops happen to be topologically linked in L . – see footnote 4 for an example.

We call a bridge *prunable* if it remains a bridge upon removal of any nontrivial subset of its loops. We say a bridge is *minimal* if it cannot be pruned. A minimal bridge of loop cardinality k is called a k -*bridge* and denoted by B_I^k . These turn out to be crucial objects for link models of entanglement to not trivialize down to hypergraphs. Within this framework, hypergraphs consists only of Hopf-linked 3-bridges. One can explicitly construct a link model from any hypergraph model by replacing vertices with infinite weight loops and (hyper)edges with loops of the corresponding weight – see Figure 1 for an illustration. The richer entanglement structure that link models capture is due to 3-bridges with different topologies as well as the existence of k -bridges with $k \geq 4$. As an example of a topologically interesting 3-bridge, consider the link model in Figure 2. If we take $I = BCD$, the loop min-cut consists of the weight-2 loop. An associated 3-bridge consists of B in the cut interior, the weight-1 loop in the cut exterior, and the weight-2 loop in the loop min-cut. This 3-bridge does not involve Hopf-linked loops – instead, as a sublink, it is a 3-component Brunnian link.

Loops in C_I can be classified by the minimal cardinality of the k -bridges for I that they contribute to. In particular, we say that $\ell \in C_I$ is a k -*loop* for I if, among all minimal bridges $B \ni \ell$ for I , the smallest bridge cardinality is k . For instance, one can easily verify that the weight-2 loop in Figure 2 is a 3-loop for $I = BCD$. We can stratify the loops $\ell \in C_I$ into subsets $C_I^{(k)} \subseteq C_I$ such that ℓ is a k -loop for I if and only if $\ell \in C_I^{(k)}$. Importantly, the subsets $C_I^{(k)}$ are disjoint for different k and define a finite spanning partition of C_I for any link with finite $|\mathcal{L}|$, i.e., $\bigcup_{k=3}^{|\mathcal{L}|} C_I^{(k)} = C_I$. When computing the total weight of a min-cut C_I , we thus have

$$|C_I| = \sum_{k=3}^{|\mathcal{L}|} |C_I^{(k)}| = \sum_{k=3}^{|\mathcal{L}|} \sum_{\ell \in C_I^{(k)}} w(\ell), \quad (3.3)$$

where $w(\ell)$ is the weight of ℓ .

3.4 A topological min-cut bottleneck

In previous work, the simple connectivity structure of graphs and hypergraphs implied that a cut, given by a subset of vertices, automatically specified the (hyper)edges whose weights contributed to the cut weight. When it comes to link models, however, loops subsume both the role of edges and vertices. While powerful for capturing a wider class of entropy vectors, this generality comes with a price. The analogous cut interior W_C now no longer specifies uniquely the loop cut C that gave rise to it and whose weight would determine the cut weight.

In order to address this new issue, one must consider the topological link classification problem [19]; that is, given a topological link, determining whether a sublink is topologically linked to the remainder of the link, or whether they are, in fact, unlinked. This problem is an outstanding problem in topology which we do not propose a solution to. Instead, we will build our formalism on top of a posited solution. We parameterize our ignorance in the form of an oracle that essentially does the difficult work of obtaining the min-cut structure of a general link.

Although such an oracle already performs a large task, we attempt to pare its power to the minimal necessary for our purposes. Specifically, consider an oracle called **Link-Min-Cut** that takes as input a link \mathcal{L} and a set of external loops $I \subseteq \partial\mathcal{L}$, and outputs its corresponding loop min-cut C_I . One can then construct the cut interior W_I as described above by taking the set of external loops I and finding the smallest number of connected sublinks that contain them after C_I is removed.

Building on the **Link-Min-Cut** routine, we employ the following additional subroutine:

$$\text{Bridge-Oracle}(\mathcal{L}, I, B) = \begin{cases} 1 & \text{if } B \text{ is a } |B|\text{-bridge across } C_I, \\ 0 & \text{otherwise.} \end{cases} \quad (3.4)$$

where B is a set of loops, a candidate minimal bridge for I , and the required knowledge of C_I is obtained from the output of **Link-Min-Cut** for the same \mathcal{L} and I arguments. The ultimate purpose of **Bridge-Oracle** is to solve a decision problem, namely, whether B serves as a $|B|$ -bridge for the set of external loops in I . In the generalization of the contraction-map proof method, this oracle will be used for determining which loops contribute to the entropy of the link model.

3.5 Generalization of contraction maps

The contraction-map method to prove entropy inequalities involves a clever partitioning of the vertex set in (hyper)graph models. For a given inequality as in (2.1) with indices $\{I_l \subseteq [n]\}_{l \in [L]}$ on its LHS, every vertex could be labeled by a bitstring of length L where each bit encoded whether it was contained in the min-cut for each I_l . For the link model, we need to generalize this addressing to loops in terms of loop min-cuts, which behave differently. Our goal in this subsection will be to prove that such a generalization exists, given access to an appropriate oracle for bridges.

We first introduce some notation and useful constructions. Pick a loop min-cut C_l for each index I_l on the LHS of the entropy inequality, and let W_l and \overline{W}_l be respectively its cut interior and exterior. To each loop $\ell \in \mathcal{L}$, we assign an upgraded bitstring $x \in \{-1, 0, 1\}^L$ via

$$x_l = \delta(\ell \in W_l) - \delta(\ell \in \overline{W}_l). \quad (3.5)$$

In other words, x_l is set to -1 , 0 or 1 if ℓ is in \overline{W}_l , C_l , or W_l , respectively. We can now partition the loop set \mathcal{L} into disjoint subsets of loops via

$$W(x) = \bigcap_{l \in [L]} W_l^{x_l}, \quad (3.6)$$

where we are declaring $W_l^1 \equiv W_l$, $W_l^0 \equiv C_l$ and $W_l^{-1} \equiv \overline{W}_l$. There will be 3^L such sets of loops (i.e., as many as bitstrings), with some possibly empty. Note that we can reconstruct the cut interior W_l of any I_l in terms of the $W(x)$ subsets via

$$W_l = \bigcup_{x: x_l=1} W(x), \quad (3.7)$$

and similarly for the loop min-cuts using

$$C_l = \bigcup_{x: x_l=0} W(x). \quad (3.8)$$

In (hyper)graph models, one was easily able to determine the edges contributing to a min-cut (cf. C_l) by just knowing the vertex min-cut (cf. W_l). In contrast, in the link model the cut interior W_l does not immediately determine what the loop min-cut C_l is in any obvious way for a general link.⁸ This explains why we needed an upgraded bitstring addressing encoding these two objects separately. In fact, we expect an explicit specification of such a map $W_l \mapsto C_l$ to be a highly nontrivial problem in knot theory which we will not try

⁸Recall that it was C_l which defined W_l , and not the converse.

to solve completely. Instead, we will try to break down this map into pieces suitable for a proof by contraction. However, with these objects defined, we can now give the formal statement of the generalized contraction map:

Proposition 1. *An entropy inequality in the form of (2.1) is valid on link models if:*

1. *There exists a map $f : \{-1, 0, 1\}^L \rightarrow \{-1, 0, 1\}^R$ from min-cuts for regions on the LHS of (2.1) to valid cuts for regions on the RHS of (2.1); that is, the regions defined by*

$$U_r = \bigcup_{x:f(x)_r=1} W(x) \quad (3.9)$$

are (not necessarily minimal) cuts for J_r .

2. *For every $3 \leq n \leq k \leq |\mathcal{L}|$, f obeys the cut-dependent contraction property*

$$\sum_{l=1}^L \alpha_l F_l^{(n,k)}(x^1, \dots, x^n) \geq \sum_{r=1}^R \beta_r \delta(f(x^1)_r = 0) \sum_{l=1}^L F_l^{(n,k)}(x^1, \dots, x^n), \quad (3.10)$$

where $F_l^{(n,k)} : \{-1, 0, 1\}^{L \times n} \rightarrow \{0, 1\}$ is an oracular indicator function that returns 1 if the input n -tuple of bitstrings corresponds to a k -bridge.⁹

We will need a few intermediate ingredients to prove this result, as we describe next. To begin, we will think of the loop cut C_l as a collection of loops independently satisfying some property with respect to W_l . Paralleling (3.3), we begin by decomposing C_l into sets of k -loops through

$$C_l = \bigcup_{k=3}^{|\mathcal{L}|} C_l^{(k)}. \quad (3.11)$$

By definition, a loop $\ell \in C_l$ is a k -loop $\ell \in C_l^{(k)}$ for the smallest k such that it belongs to a k -bridge B_l^k , where B_l^k by construction intersects all of W_l , \overline{W}_l , and C_l nontrivially. In particular, notice that if $\ell \in C_l$, then it must be that $\ell \in W(x)$ for some bitstring with $x_l = 0$. However, this decomposition is not fine-grained enough for our purposes: for a contraction condition to hold, we need to be able to sum over bitstrings, and the decomposition above does not contain enough information. Recall (3.6), which defines $W(x)$ for a bitstring x in terms of an intersections of vertex subsets W_l or their complements. However, a given k -bridge will generally intersect multiple $W(x)$, and the decomposition into k -loops does not let us determine which bridge B_l^k or set $W(x)$ it belongs to. That is, we are unable to pass directly from k -loops to a condition on bitstrings, because we cannot identify bitstrings with specific k -loops that contribute to a cut.

To handle this, we will introduce a further decomposition of a k -loop. Firstly, assuming access to an appropriate oracle for validating the existence of a bridge, we show that it is possible to decompose a cut C_l such that each loop in C_l belongs to a distinct bridge:

Lemma 1. *Consider a loop min-cut C_l and k -bridge B_l^k for the boundary region I_l . Then $|B_l^k \cap C_l| = 1$.*

Proof. We prove this by contradiction. Suppose $|B_l^k \cap C_l| > 1$, and let $\ell_i \in B_l^k \cap C_l$. By definition of bridge minimality, B_l^k is not prunable, so removing ℓ_i will cause B_l^k to no longer be a bridge. This means that $B_l^k \setminus \{\ell_i\}$ is either no longer irreducible, does not

⁹Double-counting of loop contributions is avoided by constructing $F_l^{(n,k)}$ as in Algorithm 1.

have nontrivial intersection with W_l , \overline{W}_l , and C_l , or both. However, because $\ell_i \in C_l$ and there exists at least one $\ell_j \in C_l$ with $j \neq i$, the last two options are not possible. Hence, it must be the case that $B_l^k \setminus \{\ell_i\}$ is now reducible, meaning that the sublink defined by the remaining loops contains more than one connected sublink.

If one of the resulting connected sublinks still has nontrivial intersection with W_l , \overline{W}_l , and C_l , then the resulting loop set is a k' bridge with $k' < k$, meaning B_l^k was prunable and not minimal, yielding a contradiction. If, instead, none of the resulting connected sublinks has non-trivial intersection with W_l , \overline{W}_l , and C_l , then at least one of $(B_l^k \setminus \{\ell_i\}) \cap W_l$, $(B_l^k \setminus \{\ell_i\}) \cap \overline{W}_l$, or $(B_l^k \setminus \{\ell_i\}) \cap C_l$ is topologically unlinked from the other two loop sets. If $(B_l^k \setminus \{\ell_i\}) \cap C_l$ becomes topologically unlinked from the other two loop sets, then it must be the case that, in fact, all three loop sets are topologically unlinked from each other, for otherwise C_l would have not defined a cut for the pertinent subsystem in the first place. Hence, it must be the case that $(B_l^k \setminus \{\ell_i\}) \cap W_l$ and $(B_l^k \setminus \{\ell_i\}) \cap \overline{W}_l$ are topologically unlinked, regardless of what happens to $(B_l^k \setminus \{\ell_i\}) \cap C_l$. However, this implies C_l was not minimal in weight, as we could instead define an alternative cut $C'_l = C_l \setminus \{\ell_j \in C_l : j \neq i\}$ for the same subsystem, with $(B_l^k \setminus \{\ell_i\}) \cap C_l$ now being part of either the new interior or exterior. As noted above, removing C'_l topologically unlinks the new interior and exterior, so it is a valid cut. We also see that removing C'_l makes B_l^k topologically unlinked, and the weight of C'_l is strictly smaller than that of C_l , because $C'_l \subset C_l$. We conclude that we must have $|B_l^k \cap C_l| = 1$. \square

Lemma 1 will allow us to reduce our analysis of a cut C_l to an analysis of the distinct bridges to make up each of its loops. Applying the lemma, we can separate the loops within a k -bridge B_l^k into the loop in C_l , and the remaining $k - 1$ loops in B_l^k that belong to either W_l or \overline{W}_l . Those loops in $B_l^k \cap W_l$ ($B_l^k \cap \overline{W}_l$) will necessarily belong to some $W(x)$ sets with $x_l = 1$ ($x_l = -1$). However, notice that these additional $k - 1$ loops need not all belong to distinct $W(x)$ sets – instead, in general there will be between 2 and $k - 1$ such sets containing them, as mentioned previously. We now introduce the further decomposition of loop-cuts $C_l^{(n,k)} \subseteq C_l^{(k)}$ consisting of those k -loops in C_l for which the smallest number of $W(x)$ sets needed to cover any one of its associated k -bridges is n . Enumerating bitstrings by a superscript as x^i , this means that every $\ell \in C_l^{(n,k)}$ is a k -loop with $\bigcup_{i=1}^n W(x^i)$ containing B_k minimally in the sense that n is the smallest across all associated k -bridges. Hence:

Corollary 1. *Every $C_l^{(k)}$ can be further decomposed as*

$$C_l^{(k)} = \bigcup_{n=3}^k C_l^{(n,k)} \quad (3.12)$$

where

$$C_l^{(n,k)} = \left\{ \ell \in W(x^1) : x_l^1 = 0, \bigcup_{i=1}^n W(x^i) \supseteq B_k \ni \ell \text{ minimally} \right\}_{(x^1, \dots, x^n)}, \quad (3.13)$$

where the subscript (x^1, \dots, x^n) implies iteration over all n -tuples of bitstrings in $\{-1, 0, 1\}^L$, and the choice $x_l^1 = 0$ is without loss of generality.

Notice that $C_l^{(n,k)}$ is nontrivial only if $n, k \geq 3$, for otherwise the conditions that define it are impossible to satisfy. At the level of bitstrings, we can see that the condition in (3.13) that $\bigcup_{i=1}^n W(x^i)$ contain a k -bridge minimally necessarily means that $\sum_{i=1}^n |x_l^i| = n - 1$ (i.e., only $x_l^1 = 0$), and that $|\sum_{i=1}^n x_l^i| \leq n - 3$ (i.e., there are both $+1$ and -1 bits).

While necessary, these bitstrings conditions are not sufficient to guarantee those in (3.13) – this issue will reappear when formulating the contraction map. In the contraction-map method, we want to turn a sum over the k -tuples of bitstrings satisfying the conditions in (3.13) into a sum over all k -tuples of bitstrings, with some indicator function specifying which tuples do participate in each $C_l^{(n,k)}$; however, we cannot formulate the necessary and sufficient conditions for a k -bridge purely in terms of elementary properties of a bitstring. Hence, we will need access to an oracle:

Definition 2 (Oracular Indicator). *Consider a candidate entropy inequality in the canonical form of (2.1) on an n -party link model. The oracular indicator is a Boolean map on bitstrings, $F_l^{(n,k)} : \{-1, 0, 1\}^{nL} \rightarrow \{0, 1\}$, which returns 1 when all of the necessary and sufficient conditions for a k -bridge to be minimally contained in $\bigcup_{i=1}^n W(x^i)$ are satisfied, and 0 otherwise.*

Note that, following the convention of (3.13), we may without loss of generality choose to have $F_l^{(n,k)}(x^1, \dots, x^n) = 0$ when $x_l^1 \neq 0$. The function also has one additional feature beyond checking the k -bridge conditions: it keeps track of what loops have *already* contributed to the cut weight and may be part of multiple k -bridges.¹⁰ This can be implemented by associating each loop in the cut set with the k -bridge of smallest k in which it participates, with an additional ordering for k -bridges of the same valence. An example for how such a function could be implemented, assuming access to **Link-Min-Cut** and **Bridge-Oracle** is:

Algorithm 1: Implementation of $F_l^{(n,k)}$

```

Calculate min-cut  $C_l$  for subsystem  $I_l$  via Link-Min-Cut;
Color all loops in  $C_l$  gray and set all values of  $F_l^{(n,k)}(x^1, \dots, x^n)$  to zero;
 $k := 3$ ;
while  $C_l$  contains gray loops do
   $n := 3$ ;
  while  $n \leq k$  do
    for each  $n$ -tuple of bitstrings  $(x^1, \dots, x^n)$  with  $x_l^1 = 0$  do
       $\mathbf{W} := \bigcup_{i=1}^n W(x^i)$ ;
      for each subset of  $k$  loops  $B \subseteq \mathbf{W}$  do
        if Bridge-Oracle( $\mathcal{L}, I_l, B$ ) then
           $\ell :=$  loop in  $B \cap W(x^1)$ ;
          if  $\ell$  exists and is gray then
            Color  $\ell$  green;
            Set  $F_l^{(n,k)}(x^1, \dots, x^n) = 1$ ;
          end
        end
      end
    end
     $n := n + 1$ ;
  end
   $k := k + 1$ ;
end

```

¹⁰In the context of the link given in Figure 2, notice that the weight-2 loop contributes to forming 3- but also 4-bridges, but because of the coloring strategy described below it is only counted once as a 3-loop.

To be explicit, the oracular nature of $F_l^{(n,k)}$ manifests purely in determining the sufficient conditions for B to be a k -bridge – the other features are for presentation/convenience, and can be implemented as separate functions on bitstrings.

Finally, we comment on the first condition appearing in the statement of Proposition 1, which is an extra requirement on any possible contraction map f , in addition to the contraction property. Recall that for a cut C_l for a region I_l appearing on the LHS of a candidate entropy inequality (2.1), we were able to decompose the cut as in (3.8). In order to relate the weights of these cuts to those on the RHS, we need to perform a similar decomposition of the RHS cuts. However, the caveat is that the decomposition must still be in terms of the $W(x)$ (or else the comparison would be too difficult), and the $W(x)$ are defined only in terms of the regions appearing on the LHS of (3.7). That is, given a candidate contraction map $f : \{-1, 0, 1\}^L \rightarrow \{-1, 0, 1\}^R$, when we construct the corresponding vertex subsets for the RHS as in (3.9), each such U_r must be a valid cut for the r^{th} region appearing on the RHS. This means $U_r \cap \partial\mathcal{L} = J_r$ and U_r is a cut interior for some (not necessarily minimal) loop cut C_r for J_r . If such an initial condition cannot be satisfied, then the contraction map method has simply failed. That such a function exists is immediate in the hypergraph case, due to the bijective relation between cut vertices and edges. However, it is not clear that the same is true here. We interpret this fact as simply a manifestation of another stricter condition that is required for the contraction-map method to succeed.

With these tools in hand, we now provide a proof of Proposition 1:

Proof of Proposition 1. Consider a candidate n -party entropy inequality in the canonical form of (2.1). Let C_l be the minimal loop cut for the l^{th} region appearing on the LHS. Then using the decomposition of C_l into $C_l^{(n,k)}$ from Corollary 1 and the oracular indicator $F_l^{(n,k)}$, we can write the LHS as a sum over all bitstrings:

$$\sum_{l=1}^L \alpha_l |C_l| = \sum_{l=1}^L \alpha_l \sum_{k=3}^{|L|} \sum_{n=3}^k |C_l^{(n,k)}| = \sum_{l=1}^L \alpha_l \sum_{k=3}^{|L|} \sum_{(x^1, \dots, x^n)}^{3 \leq n \leq k} F_l^{(n,k)}(x^1, \dots, x^n) |W(x^1)|. \quad (3.14)$$

We continue by considering the RHS of (2.1). By the first assumption in the statement of the proposition, there exists a function $f : \{-1, 0, 1\}^L \rightarrow \{-1, 0, 1\}^R$ such that

$$U_r = \bigcup_{x: f(x)_r=1} W(x), \quad (3.15)$$

is a valid cut for the r^{th} region appearing on the RHS. Notice that generally there are many implicit conditions f will have to satisfy. In particular, the zeros of f already specify RHS loop cuts, which themselves uniquely specify all RHS cut interiors and exteriors. This implies that all the ± 1 entries of f cannot be chosen freely, but are fixed by the zeroes of f . The latter also determine the weights of the RHS cuts U_r given by (3.9), which can be expressed as follows:

$$|U_r| = \sum_{k=3}^{|L|} \sum_{(x^1, \dots, x^n)}^{3 \leq n \leq k} \sum_{l=1}^L \delta(f(x^1)_r = 0) F_l^{(n,k)}(x^1, \dots, x^n) |W(x^1)|. \quad (3.16)$$

This expression simply counts when loops from the LHS cuts end up contributing to the RHS. Compared to (3.14), which counts loop weights in the minimum cut for each of the LHS subsystems I_l , this expression counts the weight of those loops that contribute to a cut for the RHS subsystem J_r via the assignment specified by f . Recall that the r^{th} bit

of $f(x^1)$ equaling 0 means that $W(x^1)$ is being assigned to be part of the loop cut of J_r . Altogether, for the RHS we can write

$$\sum_{r=1}^R \beta_r |U_r| = \sum_{r=1}^R \beta_r \sum_{\substack{|\mathcal{L}| \\ k=3 (x^1, \dots, x^n)}} \sum_{3 \leq n \leq k} \sum_{l=1}^L \delta(f(x^1)_r = 0) F_l^{(n,k)}(x^1, \dots, x^n) |W(x^1)|. \quad (3.17)$$

Assuming the cut specified by f as in (3.9) is indeed valid, then this will, by definition, upper bound the true value of the min-cut associated with the RHS. Hence, (2.1) would then be satisfied if

$$\begin{aligned} \sum_{\substack{|\mathcal{L}| \\ k=3 (x^1, \dots, x^n)}} \sum_{3 \leq n \leq k} \left[\sum_{l=1}^L \alpha_l F_l^{(n,k)}(x^1, \dots, x^n) \right] |W(x^1)| \\ \geq \sum_{\substack{|\mathcal{L}| \\ k=3 (x^1, \dots, x^n)}} \sum_{3 \leq n \leq k} \left[\sum_{r=1}^R \beta_r \delta(f(x^1)_r = 0) \sum_{l=1}^L F_l^{(n,k)}(x^1, \dots, x^n) \right] |W(x^1)|. \end{aligned} \quad (3.18)$$

A simple sufficient condition for this to hold is to demand that the inequality be satisfied for every n -tuple of bitstrings with $3 \leq n \leq k$ for each value of k . Were this to be the case, then all possible equations of the following form would hold:

$$\sum_{l=1}^L \alpha_l F_l^{(n,k)}(x^1, \dots, x^n) \geq \sum_{r=1}^R \beta_r \delta(f(x^1)_r = 0) \sum_{l=1}^L F_l^{(n,k)}(x^1, \dots, x^n). \quad (3.19)$$

Note that this is the second condition given in the statement of Proposition 1. If the above is true, then the given entropy inequality is also true, by the minimality property of min-cuts.¹¹ \square

We remark that we obtain a contraction condition which is *cut-dependent*, in the sense that it depends on the oracular indicator $F_l^{(n,k)}$. This cut-dependence of the contraction map is a new feature of link models that was not present in (hyper)graphs, arising from the new topological character of min-cuts.

3.6 Subadditivity and Strong Subadditivity

Although we have a generalized contraction condition in hand, an important consistency check that link models are relevant for quantum mechanics is whether the two universal quantum inequalities of subadditivity and strong subadditivity hold for topological links models. Indeed, it is known that these two inequalities completely determine the entropy cone of all quantum states for $n \leq 3$ parties.

It turns out that these properties do indeed easily follow by the definition of min-cuts on links, without requiring the contraction-map technology. In the case of subadditivity, $S(A) + S(B) \geq S(AB)$, the cuts that separate A from its complement and the cuts that separate B from its complement will by definition separate AB from its complement, and not necessarily optimally. This proves subadditivity. For strong subadditivity, $S(AB) + S(BC) \geq S(B) + S(ABC)$, the argument is much the same: the loop cuts that separate AB from its complement and BC from its complement can also be used to separate B from its complement and ABC from its complement. In addition, the subsets of loops of

¹¹Namely, since by construction for the LHS we have min-cuts with $S(I_l) = |C_l|$, whereas for the RHS we have simply constructed cuts which are not necessarily minimal and thus obey $S(J_r) \leq |C_r|$.

the AB and BC cuts that get recycled to build loop cuts for each B and ABC can be chosen to be disjoint, thus preventing double-counting on the RHS without accompanying double-counting on the LHS. The converse inequality is clearly false: loops in cuts for B and ABC do not always allow one to build a loop cut for e.g. AB , as in particular A can still be linked with C . This type of reasoning is very much the same one involved in the proof by contraction.

The above arguments show that the entropy cone of link models is contained within the n -party quantum cone for $n \leq 3$. We could show containment in the other direction by exhibiting links that realize the extreme rays of the corresponding cones. This is immediate, since any hypergraph entropies can be realized by links as explained in Figure 1, and it is known that hypergraphs realize all such extreme rays [5]. We hence find that entropy cone of link models is equivalent to the quantum entropy cone for $n \leq 3$ parties.

3.7 Potential Alternative Formulation: Super-Loops

Here, we briefly discuss a more natural generalization of the contraction-map method to loops, in the sense that it better captures how (hyper)graphs work. We are unable to prove the statement, so we leave it as a conjecture.

An intuitive way to understand the $W(x)$ sets in the graph case is as “super-vertices” – because one is only interested in the edges that enter and leave $W(x)$, the connections within a given $W(x)$ are unimportant, and one can think of each $W(x)$ as simply being a large vertex connected to other super-vertices by appropriate edges. Indeed, it is precisely this intuition that is used in the proofs of polyhedrality of the graph and hypergraph entropy cones [2, 5].

One would like to generalize this to the loop case, letting $W(x)$ denote a “super-loop”, and redefining the notion of a k -bridge to correspond to a set of k -many $W(x)$ sets. This would immediately simplify the proof and notation of the contraction condition above: we would no longer need to separately sum over n -many $W(x)$ sets to find all possible k -bridges. Instead, we would simply have to sum over all k -many subsets of the $W(x)$ sets directly.

Unfortunately, the barrier to doing so lies in, once again, the fact that there is no longer a unique way to connect loops. The key in applying the super-vertex method is that cut entropies on the new graph involving super-vertices are identical to those on the old graph. This is immediate on hypergraphs due to the simplicity of hyperedges. However, it is not clear whether the same is true for loop models. The freedom and variety of ways to connect loop sets together means that the preservation of cut entropies in the new super-loop model is not obviously guaranteed. We leave it as an open question whether such a generalization is possible.

4 Discussion

In this work, we extend the program of building models of multipartite entanglement based on min-cuts initiated in [2] for graphs and in [10, 5, 8] for hypergraphs to topological links. In particular, we motivate a specific topological generalization of the min-cut to a model of entanglement involving links where the entanglement entropy of subsystems is given by weights of linked loops. We have demonstrated the usefulness of this definition by showing that link models allow for the construction of entropy vectors which violate known inequalities for hypergraph states– see, for example, the link model in Figure 2. Additionally, we generalized the contraction-map method for proving inequalities, although this required

the inclusion of certain oracular functions. The inherent difficulty of working with the topological properties of these link models of entanglement leaves many open directions for future work.

Firstly, a more direct version of the contraction-map method with less reliance on oracles would be desirable in order to turn the proof method proposed here into a combinatorial one that can be implemented by computer. Likely this would involve formulating the necessary conditions in a cut-independent way. Alternatively, were this not to be possible, it would be interesting to prove that oracles are necessary and thus that this method just cannot be easily turned into a combinatorial one.

Relatedly, while we are able to show that link models reach beyond the hypergraph cone, we are unable to obtain insight into the geometric structure of the link cone, e.g. we do not know if the entropy cone of link models is polyhedral. It is possible that insights into the full structure of the cone may enable a more efficient or direct proof approach. We remark that lack of polyhedrality is not necessarily an issue from a conceptual standpoint: it is unknown if the quantum entropy cone is polyhedral for $n > 3$ parties, and it may well be that it is not.

Next, a heuristic strategy of constructing more “interesting” topological link models that are provably not realizable by hypergraphs would be desirable. This would require a more complete understanding of how to produce general topological links whose entropies, like those of the link in Figure 2, go beyond the limits of the hypergraph paradigm. Ultimately, one would hope to find a method for, given an entropy vector, determining whether there exists a link model realizing it and, if so, which form it takes. This would precisely correspond to a generalization of the powerful methods developed in [4] for graphs (also used in [5, 10] for hypergraphs) via integer linear programs. For link models, it is possible that nascent machine learning methods in topological link/knot classification à la those of [20] will be useful in heuristically generating these nontrivial link structures.

Moreover, even if there does exist a link model realizing an entropy vector, it may not be the only one – there may be multiple inequivalent link representations of a given entropy vector. This is similar to (hyper)graph models, where it is also possible to find arbitrarily many inequivalent graph representations of a given entropy vector. Hence, we do not believe that this non-uniqueness impacts the comparison of the link model entropy cone with the stabilizer and quantum entropy cones. It would, of course, be interesting to find “minimal” links for each entropy vector, analogous to the notion of minimality for graph representatives of entropy vectors defined by [4]. Rigorously defining a sharp notion of minimality on links would certainly be useful for distilling which properties of links characterize entanglement structures unreachable by hypergraphs. Along similar lines, it would be interesting to determine what, if any, implications the extra information contained in inequivalent representations of the same entropy vector has for physical quantum states.

Finally, and perhaps most importantly, we would like to further strengthen the connections between topological link models and physical quantum state entropies. As argued above, topological link models obey the only two inequalities known to hold on arbitrary quantum states: subadditivity and strong subadditivity (and symmetry permutations thereof). Though necessary as conditions for topological link models to capture precisely the entropies of quantum states, these are not sufficient. The following logically independent questions remain:

- *Do topological link models obey all other quantum inequalities?* If the answer to this were to be affirmative, its proof would be of paramount importance, and correspondingly a very ambitious challenge. The reason for this is that other quantum

inequalities are strongly believed to exist, but none are actually known! Obtaining an inequality valid on link models through the machinery presented in this work would shed very exciting light on this. If instead the answer were to be negative, the finding of a counterexample would be extremely interesting: it would provide a set of link entropies consistent with subadditivity and strong subadditivity, yet inconsistent with quantum mechanics, thereby pointing at the existence of novel quantum entropy inequalities that rule out such entropies.

- *Do topological link models reproduce all quantum entropies?* If true, a proof of this in the affirmative would require either an explicit construction of link models for given quantum entropy vectors, or a proof of nonexistence of valid link inequalities stronger than quantum ones, both of which seem beyond the reach of current tools. If false, however, it would only take the proof of some inequality on links that does not hold quantum mechanically, which may be more achievable using the tools in this paper.

As can be seen, the landscape of possibilities is vast and worth exploring. A YY/YN/NY/NN answer to the questions above would imply that the link entropy cone is equal to/a proper subset of/a proper superset of/incomparable to the quantum entropy cone. While the YY option would be particularly tantalizing, we believe that any of them would be of interest in advancing our tools for classifying multipartite quantum entanglement. If YN, link models would continue the program started in the context of holography in the obvious way: the most basic patterns of multipartite entanglement would be those realizable by graphs, then those realizable by hypergraphs but not graphs, then those realizable by topological links but not (hyper)graphs, etc. In fact, if NN, the same classification would still apply, since as explained around Figure 1 all hypergraph entropies are realizable by links. Of course, in this latter NN case there would however be “unphysical” link models realizing entropies which are not quantum mechanical. One may nonetheless expect that some structural constraint on link models could be implemented to avoid this, potentially bringing one back to the interesting YN case. In the NY case, an analogous constraint would also be very interesting, as it could return one to the ideal YY case.

Let us emphasize why such a classification of multipartite entanglement into graphs, hypergraphs, links, etc. is particularly exciting. Graph models of entanglement originated naturally in the context of holography [2], where the entanglement of boundary quantum states is geometrized in the bulk in a rather explicit way. This geometrization is possible precisely because of the local character of quantum correlations in holographic states. We may thus characterize quantum states having entropies realizable by graphs as possessing only *locally* entangled subsystems. Hypergraph models break this notion of locality by allowing for k -local correlations, implemented in the form of k -edges [5]. As a result, hypergraphs are further capturing quantum states which may also involve subsystems which have *non-locally* entanglement in a multi-local sense. The additional extension to links brings us quite naturally to another layer of complexity in the structural form of entanglement where multi-locality is superseded by a more global notion of *topological* entanglement. We thus see link models as an appealing extension of the classification of multipartite entanglement initiated in the context of holography in terms of graphs and followed by hypergraphs.

Given the promise demonstrated by the topological link model in going beyond the entropies describable by hypergraphs, and also given the large range of open questions for future exploration in this direction, we anticipate that this will be an open area of exploration in the area of multipartite entanglement for years to come.

Acknowledgements

We would like to thank Aaron Elersich, Greg Kuperberg, Jonathan Harper, Nils Baas, and Matthew Headrick for useful discussions. For parts of the completion of this work, N.B. was supported by the Department of Energy under grant number DE-SC0019380, and is supported by the Computational Science Initiative at Brookhaven National Laboratory, and by the U.S. Department of Energy QuantISED Quantum Telescope award. N.C. and V.P.S. were supported in part by the National Science Foundation (award number PHY-1521446), and by the U.S. Department of Energy under contract DE-AC02-05CH11231 and award DE-SC0019380. V.P.S. gratefully acknowledges support by the NSF Graduate Research Fellowship Program under Grant No. DGE 1752814. During the completion of this work, S.H.C. was supported by NSF grant PHY-1801805 and by funds from the University of California Santa Barbara, and is supported by NSF grant PHY-2107939 and by a Len DeBenedictis Graduate Fellowship.

References

- [1] Shinsei Ryu and Tadashi Takayanagi. “Holographic derivation of entanglement entropy from AdS/CFT”. *Phys. Rev. Lett.* **96**, 181602 (2006). [arXiv:hep-th/0603001](#).
- [2] Ning Bao, Sepehr Nezami, Hiroshi Ooguri, Bogdan Stoica, James Sully, and Michael Walter. “The Holographic Entropy Cone”. *JHEP* **09**, 130 (2015). [arXiv:1505.07839](#).
- [3] Sergio Hernández-Cuenca. “Holographic entropy cone for five regions”. *Phys. Rev. D* **100**, 026004 (2019). [arXiv:1903.09148](#).
- [4] David Avis and Sergio Hernández-Cuenca. “On the foundations and extremal structure of the holographic entropy cone” (2021). [arXiv:2102.07535](#).
- [5] Ning Bao, Newton Cheng, Sergio Hernández-Cuenca, and Vincent P. Su. “The Quantum Entropy Cone of Hypergraphs”. *SciPost Phys.* **9**, 067 (2020). [arXiv:2002.05317](#).
- [6] Nicholas Pippenger. “The inequalities of quantum information theory”. *IEEE Transactions on Information Theory* **49**, 773–789 (2003).
- [7] Noah Linden, František Matúš, Mary Beth Ruskai, and Andreas Winter. “The Quantum Entropy Cone of Stabiliser States”. *LIPICs* **22**, 270–284 (2013). [arXiv:1302.5453](#).
- [8] Michael Walter and Freek Witteveen. “Hypergraph min-cuts from quantum entropies”. *J. Math. Phys.* **62**, 092203 (2021). [arXiv:2002.12397](#).
- [9] Sepehr Nezami and Michael Walter. “Multipartite Entanglement in Stabilizer Tensor Networks”. *Phys. Rev. Lett.* **125**, 241602 (2020). [arXiv:1608.02595](#).
- [10] Ning Bao, Newton Cheng, Sergio Hernández-Cuenca, and Vincent Paul Su. “A Gap Between the Hypergraph and Stabilizer Entropy Cones” (2020). [arXiv:2006.16292](#).
- [11] Grant Salton, Brian Swingle, and Michael Walter. “Entanglement from Topology in Chern-Simons Theory”. *Phys. Rev. D* **95**, 105007 (2017). [arXiv:1611.01516](#).
- [12] Vijay Balasubramanian, Jackson R. Fliss, Robert G. Leigh, and Onkar Parrikar. “Multi-Boundary Entanglement in Chern-Simons Theory and Link Invariants”. *JHEP* **04**, 061 (2017). [arXiv:1611.05460](#).
- [13] Sungbong Chun and Ning Bao. “Entanglement entropy from SU(2) Chern-Simons theory and symmetric webs” (2017). [arXiv:1707.03525](#).
- [14] Sergey Mironov. “Topological Entanglement and Knots”. *Universe* **5**, 60 (2019).

- [15] Louis H. Kauffman and Eshan Mehrotra. “Topological aspects of quantum entanglement”. *Quantum Inf Process***18** (2019).
- [16] D. Aharonov, V. Jones, and Zeph Landau. “A polynomial quantum algorithm for approximating the jones polynomial”. *Algorithmica* **55**, 395–421 (2006).
- [17] Chris Akers, Sergio Hernández-Cuenca, and Pratik Rath. “Quantum Extremal Surfaces and the Holographic Entropy Cone”. *JHEP* **11**, 177 (2021). [arXiv:2108.07280](#).
- [18] M Hein, Jens Eisert, and Hans Briegel. “Multiparty entanglement in graph states”. *Phys. Rev. A* **69**, 062311 (2004).
- [19] Nathan Habegger and Xiao-Song Lin. “The classification of links up to link-homotopy”. *Journal of the American Mathematical Society* Pages 389–419 (1990).
- [20] Sergei Gukov, James Halverson, Fabian Ruehle, and Piotr Sułkowski. “Learning to Unknot”. *Mach. Learn. Sci. Tech.* **2**, 025035 (2021). [arXiv:2010.16263](#).