Fault-Tolerant Logical Gates in Holographic Stabilizer Codes Are Severely Restricted

Sam Cree^{1,*}, Kfir Dolev¹, Vladimir Calvera², and Dominic J. Williamson¹

¹*Stanford Institute for Theoretical Physics, Stanford University, Stanford, California 94305, USA*

²*Department of Physics, Stanford University, Stanford, California 94305, USA*

(Received 23 April 2021; accepted 12 July 2021; published 31 August 2021)

We evaluate the usefulness of holographic stabilizer codes for fault-tolerant quantum computing by studying their allowed sets of fault-tolerantly implementable gates. We treat them as subsystem codes and show that the set of transversally implementable logical operations is contained in the Clifford group for sufficiently localized logical subsystems. As well as proving this concretely for several specific codes, we argue that this restriction naturally arises in any stabilizer subsystem code that comes close to capturing certain properties of holography. We extend these results to approximate encodings, locality-preserving gates, and certain codes whose logical algebras have nontrivial centers, and discuss cases where restrictions can be made to other levels of the Clifford hierarchy. A few auxiliary results may also be of interest, including a general definition of an entanglement wedge map for any subsystem code, and a thorough classification of different correctability properties for regions in a subsystem code.

DOI: [10.1103/PRXQuantum.2.030337](https://doi.org/10.1103/PRXQuantum.2.030337)

I. INTRODUCTION

The anti-de Sitter/conformal field theory (AdS/CFT) correspondence [1,2] is a particularly fruitful instance of the holographic principle [3,4] from high-energy theory—it refers to a duality between a theory of quantum gravity in $D + 1$ dimensions and a conformally invariant quantum field theory on the D -dimensional boundary of the original space. It was observed in Ref. [5] that this duality has many properties in common with quantum error-correcting codes, which motivated the development of several families of error-correcting codes built out of tensor networks [6–14] as toy models for the holographic correspondence. These *holographic codes* encode information into a series of degrees of freedom that are increasingly nonlocal, with the degree of nonlocality manifesting itself in the logical system as an additional emergent geometric dimension. The logical system is associated with this higher-dimensional “bulk” geometry, and the physical system is associated with the lower-dimensional “boundary.”

Quantum error correction is a remarkable idea that unlocks the possibility of feasible quantum computation

[15–17]. Many of the leading proposals for implementing quantum error correction are based on a concept borrowed from theoretical condensed matter physics—namely, topological order [18–20]. These *topological codes* [21–23] encode information into the (necessarily nonlocal) degrees of freedom that are sensitive to the topology of the surface on which the code lives. Thus, they protect against local, topologically trivial, operations. The success of topological codes motivates us to assess the usefulness of these new holographic codes for practical quantum computation.

There are a number of reasons to suspect that holographic codes may be of practical use for quantum computing. Holographic codes can admit erasure thresholds comparable to that of the widely studied surface code [6,11], and likewise for their threshold against Pauli errors [12]. Their holographic structure also naturally leads to an organization of encoded qubits into a hierarchy of levels of protection from errors, which could be useful for applications that call for many qubits with varying levels of protection. In particular, this is reminiscent of many schemes for magic state distillation [24,25]—and indeed, the concatenated codes [26] used for magic state distillation have a hierarchical structure similar to that of holographic codes. The layered structure of holographic codes is also reminiscent of memory architectures in classical computers, where it is useful to have different levels of short-term and long-term memory. Although these codes have some notable drawbacks, in particular holographic stabilizer codes require nonlocal stabilizer generators, other codes, such as concatenated codes, suffer similar

*scree@stanford.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

drawbacks and have still proven to be useful. Conversely, the stringent requirement of nonlocal stabilizer generators allows holographic codes to protect many more qubits than a topological code and attain a finite nonzero encoding rate, which is typically not possible for topological codes. Nonetheless, many open questions remain about the usefulness of holographic codes for fault-tolerant quantum computing.

Beyond the robust storage of quantum information, a desirable feature of a quantum codes is the ability to perform a large set (ideally a universal set) of logical operations while maintaining the protection provided by the code. Thus, it is useful to study the set of *fault-tolerant* logical gates implementable in a given code (i.e., those that confine the spread of a physically relevant set of correctable errors sufficiently well that they remain correctable). The simplest examples are *transversal* gates, which do not couple the physical subsystems of the code and hence do not spread errors at all. The desirable error correction properties of these gates come at a cost, which is that they are restricted to implementing a nonuniversal (and finite) group of logical operations on a code with nontrivial distance [27].

Topological codes possess additional geometric structure, which results in a wider class of fault-tolerant gates that subsumes the set of transversal gates: those implementable by *locality-preserving* operations. Locality-preserving operations are those which, when acting on a local operator via conjugation, expand the region on which it acts by a bounded amount. These preserve the correctability of errors confined to sufficiently small regions. However, like transversal gates, locality-preserving gates in topological codes are also restricted to form a finite group [28–30]. Specifically, it has been shown that locality-preserving gates in topological *stabilizer subsystem codes* [28,31,32] in D spatial dimensions can implement only gates from $\mathcal{C}_D$, the D th level of the Clifford hierarchy, which is a series of increasing sets of gates that includes the Pauli group ($\mathcal{C}_1$) and the Clifford group ($\mathcal{C}_2$). These sets obey $\mathcal{C}_1 \subset \mathcal{C}_2 \subset \dots$, and the bound is known to be saturated; i.e., for all $D > 1$, one can find a D -dimensional topological stabilizer (subsystem) code that can transversally implement a gate from $\mathcal{C}_D \setminus \mathcal{C}_{D-1}$, for example, the color code [33,34].

Just as for topological codes, the geometric structure of holographic codes endows locality-preserving gates with fault tolerance. This raises a natural question: what restrictions does the holographic structure of these codes impose on locality-preserving gates? A natural guess for holographic *stabilizer* codes in D spatial dimensions might be that one can access $\mathcal{C}_{D+1}$, the $(D+1)$ th level of the Clifford hierarchy, due to their locality in the extra holographic dimension and the analogy to topological codes. Here we show that the restrictions are much severer than this—namely, that locality-preserving gates for such

holographic codes are dimension-independently restricted to the Clifford group, $\mathcal{C}_2$.

A. Summary of results and main ideas

In this work, we show that for sufficiently locality-preserving physical gates, the restricted logical action on a sufficiently small region of the bulk cannot implement a non-Clifford gate for any stabilizer code with a holographic structure that captures certain properties of AdS/CFT in any spatial dimension. Deeper into the bulk, this restriction applies even to unitaries that expand the support of local operators by increasingly large amounts. We describe a number of generalizations of our result, including an extension to approximate stabilizer encoding maps and to the codes studied in Ref. [8] with logical algebras whose centers are nontrivial; in all these cases we show that it remains impossible to implement a logical non-Clifford gate. We also show that as our assumptions are relaxed, then even if non-Clifford gates are allowed, the logical gate is still restricted to some higher level of the Clifford hierarchy determined by the size of the logical bulk region and the degree of spreading by the locality-preserving unitary.

We now explain the relatively simple idea that underlies most of our results. First, we recall the following implication of a lemma due to Pastawski and Yoshida [31] stated informally: any stabilizer subsystem code whose physical qudits can be partitioned into three correctable regions can only support transversal operators that implement Clifford gates (see Lemma II.18 for the technical statement). For any given instance of a holographic code and choice of a small bulk subregion, it is typically easy to find a partitioning of the physical system into three correctable regions, meaning this lemma can be applied. The main contribution of this work is to argue from the physical principles of holography when and why it is possible to find such a tripartition. A key error-correcting property of holographic codes is a relationship between the correctability of two complementary regions, known as complementary recovery. We argue that so long as regions that *almost* satisfy complementary recovery are not *too* rare, then the partitioning into three correctable regions is possible; see Fig. 1. We go on to make use of the geometric structure and locality inherent to the physical qubits at the boundary of a holographic code to extend our results to cover locality-preserving unitaries. Informally, we show the following.

Theorem I.1 (informal statement of Theorem IV.6). *If there exists an appropriate region that almost satisfies complementary recovery and the protected bulk subsystem is sufficiently small, then any sufficiently locality-preserving unitary with well-defined logical action implements an element of the Clifford group.*

FIG. 1. A partition of the physical qubits on the boundary into three regions, R_1 , R_2 , and $(R_1 \cup R_2)^c$, that are all correctable with respect to some logical qubits associated with a bulk region. This is argued in detail later in the paper using the holographic structure of the code; in particular properties of the entanglement wedge of $R_1 \cup R_2$ depicted in green and labelled EW. In this setting, locality-preserving logical gates U that spread local operators only by an amount S_U can implement only Clifford gates on the logical qubits.

There already exists a bound [31] establishing that transversal unitaries acting on logical subsystems that have sufficiently large distances, such as those associated with qubits deep in the bulk, are already confined to a particular level of the Clifford hierarchy. Our work can be seen as a strengthening of that result in the specific context of holographic codes. We apply it to nontransversal but locality-preserving gates, as well as logical qubits associated with shallower regions of the bulk whose distances are not sufficiently large for the previously established arguments to apply.

Our results share some similarities with recent work ruling out global symmetries in AdS/CFT [35,36]. The results proved there imply that holographic codes satisfying a set of assumptions motivated by the AdS/CFT correspondence do not support any nontrivial locality-preserving gates that act uniformly on the bulk encoded qubits. Our results differ on one hand as we require weaker assumptions to be satisfied by the holographic codes in question, and we consider a broader class of locality-preserving gates that may act on the bulk in a nonuniform fashion. As a consequence, our results are in a sense weaker where they overlap, ruling out only non-Clifford global symmetries in the bulk rather than all global symmetries. Our broader assumptions necessarily require this weakening of the results, as our notion of

holographic codes includes the well-known HaPPY code [31], which is known to violate the stronger “no global symmetries” result by admitting nontrivial Clifford global symmetries [37]. On the other hand, we also assume additional structure by requiring the codes to be stabilizer codes, which allows us to make stronger claims about bulk logical operations that are not global symmetries.

As far as practical applications are concerned, unfortunately our main results point out an obstacle for any potential application of holographic stabilizer codes for magic state distillation. Specifically, we show that such codes essentially do not support transversal logical gates outside the Clifford group as required for magic state distillation.

B. Outline

This article is organized as follows: We first introduce the background on subsystem codes in Sec. II and prove some of their properties that are useful for our purposes. We also introduce a general definition of entanglement wedge, termed the “maximal entanglement wedge,” which is defined for an arbitrary subsystem code without any assumptions of additional structure such as a particular geometry. This contributes to a growing set of such results [38–40] that identify holographic features that emerge from surprisingly simple quantum error-correction properties. We then describe some of the properties that we associate with holographic codes, such as complementary recovery and a geometric structure. We also review the stabilizer formalism to emphasize some crucial subtleties that allow a natural definition of the Clifford group on the logical space. Finally, we review the background on fault-tolerant implementation of logical gates; in particular, a crucial lemma from Ref. [31] (the Pastawski-Yoshida lemma) that shows that the existence of regions with certain error-correction properties implies a restriction on the logical gates implementable by locality-preserving operations.

In Sec. III, we consider a prototypical family of holographic codes constructed from perfect tensors, known as HaPPY codes after the authors of Ref. [6]. We show that any code made from copies of a single perfect tensor must support appropriate regions to apply the Pastawski-Yoshida lemma, and we argue that this property extends to any HaPPY code built from more general combinations of perfect and planar-perfect tensors.

Since the introduction of HaPPY, many other families of holographic stabilizer codes have been introduced that extend beyond the assumptions made in Ref. [6]. Thus, we argue in Sec. IV on more general grounds that any holographic code that even almost captures the code properties of AdS/CFT (*complementary recovery* in particular) must also have their locality-preserving gates restricted to

implement only Clifford gates. We also extend this argument to approximate codes such as that in Ref. [9], and to codes whose logical algebras may have nontrivial centers [8]. Finally, in Sec. IV G we show variants of this argument for other levels of the Clifford hierarchy. For example, even if a bulk region is sufficiently large that the arguments of Sec. IV do not apply, we can still establish a weaker restriction to a level $\mathcal{C}_n$ for some $n > 2$. In other cases we are able to show an even stronger restriction than our main result. In these cases, an operator that acts only on the logical subsystem and trivially elsewhere cannot even implement a Clifford gate unless it is a Pauli operator (i.e., it is restricted to $\mathcal{C}_1$).

II. PRELIMINARIES

In this section, we introduce various properties and additional structures of error-correcting codes that are necessary to formulate and prove the main results of our work in later sections. We review subsystem and stabilizer codes, introduce the error-correction and geometric properties of holographic codes, and review some useful lemmas about locality-preserving gates.

In holographic codes, the bulk space (which we denote B) is associated with the encoded space of the code, and the boundary space (denoted ∂B) is associated with the physical degrees of freedom of the code. Different encoded/bulk qubits have different levels of protection against error. Those near the boundary are susceptible to small errors occurring nearby on the boundary, whereas those deep in the bulk are well protected against all small errors. For this reason, it is useful to characterize the error-correction properties of such a code with respect to a choice of a subset of the encoded qubits that are intended to be protected while the remainder need not be. This can be accomplished with the subsystem code formalism, which we introduce here.

A. Subsystem codes

For a general error-correcting code, the physical Hilbert space $\mathcal{H}_p$ decomposes as $\mathcal{H}_p = \mathcal{H}_c \oplus \bar{\mathcal{H}}_c$, with $\mathcal{H}_c$ the code subspace and $\bar{\mathcal{H}}_c$ its complement. For a standard code, the code subspace is taken to be isomorphic to a logical system $\mathcal{H}_L$ that represents the important encoded quantum information. In the more general framework of *subsystem* codes, $\mathcal{H}_c$ is isomorphic to a tensor product $\mathcal{H}_L \otimes \mathcal{H}_J$, with the logical subsystem $\mathcal{H}_L$ storing the “important information,” and the *junk* subsystem (sometimes known as the *gauge* subsystem) [41] $\mathcal{H}_J$ storing “unimportant information.” This division is particularly useful when different subsystems of the codespace experience different levels of protection against error, as one can evaluate the error-correction properties of $\mathcal{H}_L$ independently of $\mathcal{H}_J$, which may be highly susceptible to errors.

We sometimes refer to the combined system $\mathcal{H}_L \otimes \mathcal{H}_J$ as the “encoded space,” $\mathcal{H}_e = \mathcal{H}_L \otimes \mathcal{H}_J \cong \mathcal{H}_c$.

The encoding isometry $V: \mathcal{H}_L \otimes \mathcal{H}_J \rightarrow \mathcal{H}_p$ maps the combination of these two subsystems into the physical Hilbert space. Its image is the code subspace, $\mathcal{H}_c$, and we denote the projector onto this subspace as $\Pi = VV^\dagger \in \mathcal{B}(\mathcal{H})$, where $\mathcal{B}(\mathcal{H})$ denotes the set of operators acting on the Hilbert space $\mathcal{H}$. An operator $A \in \mathcal{B}(\mathcal{H})$ is said to be codespace-preserving (CSP) if it commutes with the code projector, $[A, \Pi] = 0$. As we only aim to protect the logical subsystem from error, a codespace-preserving operator that entangles the two encoded subsystems is undesirable, as it can contaminate the important information with junk information that has not been well protected from error. Thus, we focus on *dressed CSP operators*, which are CSP operators A that act as a product operator on the two subsystems, i.e., $V^\dagger A V = A_L \otimes A_J \in \mathcal{B}(\mathcal{H}_L \otimes \mathcal{H}_J)$. We say that such an operator A *implements* the logical operation $A_L \otimes A_J$, which is a *dressed logical operator*. When the operator acts trivially on the junk subsystem, i.e., $A_J = \mathbb{1}_J$, we say that A is a *bare CSP operator*, and its implemented operator $A_L \otimes \mathbb{1}_J$ is a *bare logical operator*. We say that a pair of bare CSP or dressed CSP operators A and B are *logically equivalent* if, up to scalar multiplication, they act the same way on the logical subsystem (i.e., $A_L \propto B_L$).

We now discuss error-correction properties of subsystem codes. We present three equivalent conditions for a particular region R (i.e., a collection of physical subsystems) to be *correctable* [42]—i.e., when it is possible to correct for the erasure of R . The conditions are expressed in terms of operators that are supported in R , by which we mean those elements of $\mathcal{B}(\mathcal{H})$ that take the form $A_R \otimes \mathbb{1}_{R^c}$.

Lemma II.1. *For a region R of a physical Hilbert space of a code with projector Π , the following properties are equivalent [31]:*

- (a) *There exists a recovery channel $\mathcal{R}$ such that for every bare CSP operator A , $\Pi(\mathcal{R} \circ \mathcal{N}_R)^\dagger(A)\Pi = \Pi A \Pi$, with $\mathcal{N}_R(\cdot) := \text{tr}_R(\cdot)$ the channel that completely erases subsystem R .*
- (b) *For any operator B_R supported in R and any bare CSP operator A , $[A, \Pi B_R \Pi] = 0$.*
- (c) *For any logical operator A_L , there exists a bare CSP operator A supported in R^c such that A implements $A_L \otimes \mathbb{1}_J$.*

Definition II.2. *A region R of the physical Hilbert space is said to be correctable if the above conditions hold.*

Each of these conditions refers to bare CSP operators, which is in turn dependent on the choice of logical subsystem $\mathcal{H}_L$. Some regions may be correctable for some choices of $\mathcal{H}_L$ and not others. In the case that the junk subsystem is trivial, i.e., $\mathcal{H}_L \equiv \mathcal{H}_c$, then the equivalence of the

first two points in Lemma II.1 corresponds to the familiar Knill-Laflamme error-correction conditions [24], and the equivalence of the third point is the so-called *cleaning lemma* for stabilizer codes [43].

In standard codes, there is another property equivalent to the three presented in Lemma II.1, which is that any CSP operator A supported in R implements the logical identity up to a scalar, $V^\dagger A V \propto \mathbb{1}$. A similar property holds for correctable regions in subsystem codes, although it is not known to be a sufficient criterion for correctability in this context.

Lemma II.3. *Suppose a region R is correctable. Then any dressed CSP operator supported on R implements a logical operation of the form $\mathbb{1}_L \otimes A_J$ for some $A_J \in \mathcal{B}(\mathcal{H}_J)$.*

Proof. Suppose R is correctable. Then take any dressed CSP operator A supported on R with $V^\dagger A V = A_L \otimes A_J$. By assumption, for any $B_L \in \mathcal{B}(\mathcal{H}_L)$, one can find B supported on R^c such that $V^\dagger B V = B_L \otimes \mathbb{1}_J$. Since A and B are supported on complementary regions, we have $[A, B] = 0$, which again implies $[A_L, B_L] = 0$. Since this is true for any $B_L \in \mathcal{B}(\mathcal{H}_L)$, A_L must be proportional to the identity, $A_L = c\mathbb{1}_L$. ■

B. Entanglement wedge in subsystem codes

The encoded space of a subsystem code factorizes into a tensor product of a logical subsystem and a junk subsystem. More generally, it typically has a richer tensor product structure in the form of a decomposition into local encoded subsystems

$$\mathcal{H}_e = \bigotimes_{i=1}^{n_L} \mathcal{H}_i, \quad (1)$$

for example, individual qudits or sites of a lattice. Then each subset $S \subset \{1, \dots, n_L\}$ corresponds to a choice of logical subsystem $\mathcal{H}_L = \bigotimes_{i \in S} \mathcal{H}_i$, with junk subsystem $\mathcal{H}_J = \bigotimes_{i \notin S} \mathcal{H}_i$. We then say that a region is correctable with respect to S if it is correctable with respect to the corresponding choice of logical and junk subsystems.

The correctability of a physical region is preserved under union of logical subsystems, i.e., if a region R is correctable with respect to S_1 and S_2 , it is also correctable with respect to $S_1 \cup S_2$ [44]. This implies the existence of a unique “maximal” logical subsystem given by the union of all local encoded subsystems with respect to which R is correctable. We define this to be the *maximal entanglement wedge* of R^c in the subsystem code.

Definition II.4. *The maximal entanglement wedge, denoted $\mathcal{E}[R]$, is defined to be the union of all local*

encoded subsystems with respect to which R^c is correctable.

We justify the terminology used here by comparison with the entanglement wedge as defined in holography. Any operator acting on local encoded subsystems in the entanglement wedge of R can be reconstructed in R , just as in holography. Furthermore, any alternative definition of an entanglement wedge with this property (e.g., the greedy entanglement wedge map of HaPPY), must be included within the maximal entanglement wedge as defined above [45].

This definition of the maximal entanglement wedge makes no additional assumptions beyond a particular choice of decomposition into local encoded subsystems. Thus, the existence of an entanglement wedge map is a general feature of any subsystem code.

Relevant to our purposes of assessing the usefulness of holographic codes for fault-tolerant quantum computation, the maximal entanglement wedge suggests a correlation between the usefulness of a subsystem code for practical error correction against erasure of R^c , and its quality as a model of holography—namely, both of these characteristics are improved when the entanglement wedge of R is larger. The former should be clear; a larger entanglement wedge means that R^c is correctable with respect to more subsystems, and so more encoded information can be protected from erasures.

For the latter, there is a limit to how large entanglement wedges can be—a limit that, it turns out, AdS/CFT saturates. A given local encoded subsystem cannot belong to both the maximal entanglement wedge of a boundary region R and its complement R^c , as this would imply that its bare logical operators can be reconstructed in both R and R^c , which would mean that all its bare logical operators must commute; a contradiction. Thus, we have, in general,

$$\mathcal{E}[R] \subseteq \mathcal{E}[R^c]^c. \quad (2)$$

If this limit is saturated for some R , i.e., $\mathcal{E}[R] = \mathcal{E}[R^c]^c$, then we say that R obeys *complementary recovery*.

Definition II.5. *If all local encoded subsystems belong to either $\mathcal{E}[R]$ or $\mathcal{E}[R^c]$, then R is said to obey complementary recovery.*

In full AdS/CFT, the entanglement wedge map is expected to satisfy this property for any region R —in this case we say that the code itself obeys complementary recovery. In fact, it is known that a subsystem code with complementary recovery automatically exhibits the Ryu-Takayanagi formula from holography [38]. Thus, codes with larger entanglement wedges—closer to the limit set

by Eq. (2)—are both better at protecting from errors [46] and more closely resemble holography.

Following this line of reasoning, several previous studies defined holographic codes to be those that satisfy complementary recovery. As we show later in Theorem IV.1, our results are particularly straightforward to state for these codes. However, this class of codes, which we refer to as *ideal holographic codes*, is too restrictive for our purposes. In particular, the most notable (and original) example of a holographic stabilizer code, the HaPPY code, is excluded from this class. It is relatively easy in HaPPY to construct a region with many disconnected pieces such that most of the bulk qubits lie in neither its entanglement wedge nor that of its complement—an example is shown in Fig. 2(c) where both entanglement wedges are completely empty. One can also construct connected regions in HaPPY codes that do not obey complementary recovery [see, e.g., Fig. 2(b)]; however, these typically have only a small fraction of the bulk that lies in neither entanglement wedge. Furthermore, many regions do not have any such residual bulk region, as in Fig. 2(a). Other holographic codes with promising attributes for practical application similarly fail to meet the stringent criterion of an ideal holographic code [11,12]. Thus, for our purposes, a holographic code is one which (*almost*) saturates Eq. (2) for many regions R ; the more common this property is among the code's regions, the closer to an ideal holographic code it becomes. For example, connected regions in HaPPY typically fail to satisfy complementary recovery by a small fraction of the total number of bulk sites, tending to zero in the worst case [47] as $\log |R|/|B| \rightarrow 0$ as the number of sites in the boundary diverges. Hence, the connected regions in HaPPY *almost* satisfy complementary recovery, making it a close model of holographic features for connected boundary regions.

C. Geometric structure of holographic codes

To this point, we have relied only on the tensor product structure of the physical and logical spaces, and not on the associated geometric structures that are typically associated with holographic codes (and assumed in pictures such as in Fig. 2). Although some of our results hold without a geometric structure, we require geometric structure at times—for example, to generalize beyond transversal gates to locality-preserving gates in Sec. IV E.

Following Ref. [39], each local encoded subsystem in the tensor product decomposition of Eq. (1) is associated with a specific point in some continuous bulk geometry. A geometric region is then associated with the tensor product of Hilbert spaces corresponding to each local encoded subsystem located in that region. We typically identify geometric distances in the bulk with the graph distances along the graph that underlies the tensor network of the

FIG. 2. Assorted examples where complementary recovery holds or fails in HaPPY. This and following figures are adapted from figures in Ref. [6]. (a) An example of a connected region R in a HaPPY code that obeys complementary recovery; i.e., its entanglement wedge is the complement of the entanglement wedge of its complement. The white nodes in the yellow shaded region denote R , and the red nodes in the yellow shaded region denote its entanglement wedge $\mathcal{E}[R]$. The corresponding nodes in the red shaded region denote R^c and $\mathcal{E}[R^c]$, respectively. We use this coloring and notation throughout Figs. 2 and 3. (b) A case in which a connected region does not obey exact complementary recovery, and instead a small $\mathcal{O}(1)$ residual region exists outside both entanglement wedges. (c) An example of a choice of a highly disconnected region R such that both it and its complement have empty entanglement wedges. This is an extreme case in which complementary recovery maximally fails to hold.

encoding isometry. The boundary geometry is then inherited by smoothly continuing this bulk geometry to the outward-facing legs of the outermost layer of tensors. We use the geometric structure in two ways; first, we sometimes refer to *connected* regions of the boundary, which are defined in terms of this geometry (e.g., in the above discussion about connected regions in HaPPY). Second, we occasionally use distance measures in either the bulk or the boundary; these are defined in terms of the graph

distance of the relevant lattice. We assume these distance measures are normalized such that two adjacent sites are a distance of order 1 apart.

We now introduce some characterizations of how well protected a logical subsystem S is against errors, some of which make explicit use of the geometric structure associated with holographic codes.

Definition II.6. We define the code distance $d(S)$ of a logical subsystem S in a subsystem code as the size of the smallest boundary region R that is not correctable with respect to S , with size measured by the number of qubits. This is equivalent to

$$d(S) = \min_R \{|R| : R \text{ is not correctable with respect to } S\}. \quad (3)$$

Physically motivated topological and holographic codes come endowed with a richer geometric structure. This leads to refined notions of generalized code distance, several of which we now introduce with holographic codes in mind.

Definition II.7. We define the connected code distance $d_c(S)$ of a logical subsystem S of a subsystem code as the size of the smallest connected boundary region R that is noncorrectable with respect to S , with size measured by the number of qubits. This definition satisfies $d_c \geq d$.

Finally, we may wish to incorporate the geometric structure in any number of other ways. We define a notion of distance that depends on an arbitrary “size” function, $\sigma(R)$, defined for any region R . Typically, a useful definition will satisfy other constraints, such as $R' \subset R \implies \sigma(R') < \sigma(R)$ or $\sigma(R) \geq 0$, but we do not require it.

Definition II.8. Take any notion of “size” of a region, $\sigma(R)$. We define the σ -code distance $d_\sigma(S)$ of a logical subsystem S of a subsystem code as the minimum size of a noncorrectable region, i.e.,

$$d_\sigma(S) = \min_R \{\sigma(R) : R \text{ is not correctable with respect to } S\}. \quad (4)$$

Definition II.7 can be thought of as a special case in which σ is defined to be infinity for nonconnected regions and to be the number of qubits in the region otherwise.

Along with the above definitions of distance, we recall the notion of *code price* [39], defined as follows.

Definition II.9. We define the code price $p(S)$ of a logical subsystem S of a subsystem code as the minimum size of a

region that can reconstruct all logical operators on S , with size measured by the number of qubits. This is equivalent to

$$p(S) = \min_R \{|R| : R^c \text{ is correctable with respect to } S\}. \quad (5)$$

As with distance, equivalent versions of the connected code price p_c and the σ -code price p_σ can be defined. The following is a straightforward corollary of Eq. (2).

Lemma II.10. The price is larger than the distance, $p(S) \geq d(S)$.

We expect the entanglement wedge map in a holographic stabilizer code to be compatible with its geometric structure, as is the case in AdS/CFT. Specifically, we expect that a local extension of the bulk region should be possible by locally deforming the boundary region—the entanglement wedge map is “smooth” in some sense [48]. Of course, the hyperbolic nature of geometries associated with spatial slices of AdS means that an $\mathcal{O}(1)$ extension of the entanglement wedge in the bulk by an amount δ may require an $\mathcal{O}[\exp(r + \delta)]$ extension to the diameter of the boundary region, where r is the distance to the boundary in AdS units (the continuous analog of the number of layers in HaPPY). Nevertheless, this weak notion of continuity is sufficient for our purposes, as the distance of a bulk region typically also scales exponentially with r . The following definition formalizes this notion.

Definition II.11. A code is said to have a κ_R -smooth entanglement wedge map if for any choice of a connected boundary region R and positive numbers δ and x that satisfy $\kappa_R(\delta) \leq x$ there exists a connected boundary region R_x such that $R \subset R_x$, $|R_x - R| \leq x$ and $B[\mathcal{E}(R), \delta] \subset \mathcal{E}(R_x)$, where $B[\mathcal{E}(R), \delta]$ is the region of all points with distance at most δ from $\mathcal{E}[R]$. In most cases we take κ_R to be a nondecreasing function that satisfies $\kappa_R(0) = 0$.

In the above definition, κ_R measures how much a boundary region R must be extended to accommodate a small extension of its entanglement wedge $\mathcal{E}(R)$ in the bulk. This is portrayed in Fig. 3.

D. Stabilizer subsystem codes

The main results of this paper refer to the Clifford hierarchy, which is closely related to the Pauli group that arises in stabilizer codes. In this section we carefully introduce this notion to point out a few subtleties.

Given a decomposition of a Hilbert space into qudits and a particular choice of basis for each qudit, one can construct a Pauli group as follows. The Hilbert space decomposes into d -dimensional qudits, $\mathcal{H} = \bigotimes_{n=1}^N \mathcal{H}_n$, with a specific basis $\{|j\rangle_n\}_{j=1}^d$ chosen for each $\mathcal{H}_n \cong \mathbb{C}^d$. The

FIG. 3. Illustration of the idea of a κ_R -smooth entanglement wedge map, as defined in Definition II.11. (a) An example of a region R with small κ_R . The region R is shown in red, and for an expansion of one graph distance in the bulk, a boundary expansion of only one site in either direction is required; thus, $\kappa_R(1) = 1$. The expanded region is shown in yellow. (b) An example of a region R with a large κ_R . Expansion by 1 in the bulk requires at least 11 sites added to either side; thus, $\kappa_R(1) = 11$.

Pauli group $\mathcal{P}$ is generated by generalized X operators $X_n = \sum_{j=1}^d |(j+1) \bmod d\rangle_n \langle j|$ (with identities acting on all but the n th qudit) and generalized Z operators $Z_n = \sum_{j=1}^d \omega_d^j |j\rangle_n \langle j|$ (with ω_d a primitive d th root of unity).

A stabilizer code requires the physical space to be equipped with a specific tensor product structure and choice of local basis, and some Abelian subgroup of the Pauli group $\mathcal{S} \subset \mathcal{P}$ such that the only element of $\mathcal{S}$ in the center [49] of $\mathcal{P}$ is the identity, $\mathcal{Z}[\mathcal{P}] \cap \mathcal{S} = \{1\}$ [50]; this group $\mathcal{S}$ is called the “stabilizer group.” The code subspace $\mathcal{H}_c$ is then characterized as the space of states that are +1 eigenstates of all of elements $S \in \mathcal{S}$. For a given choice of stabilizer group, there exists a unique [51] tensor product structure and local basis for the encoded space $\mathcal{H}_e \cong \mathcal{H}_c$ such that the codespace-preserving physical Pauli operators implement elements of the encoded Pauli group [52]. A stabilizer subsystem code corresponds to a particular division of these encoded qudits into logical and junk subsystems such that the logical subsystem has at least one qudit [53].

When the junk subsystem is not trivial, it may be possible that no operator supported on R^c can implement the logical Pauli operator $A_L \otimes 1_J$, but some dressed CSP operator can implement $A_L \otimes A_J$. Thus, it is useful to

introduce a weaker correctability property that we term “dressed-cleanable,” following Ref. [31].

Definition II.12. R is dressed-cleanable if for any logical Pauli operator P_L there exists a dressed CSP operator supported on R^c that implements $P_L \otimes P_J$ for some operator $P_J \in \mathcal{B}(\mathcal{H}_J)$.

This is a strictly weaker requirement than correctability. Compare both notions using tensor networks in Fig. 4.

Lemma II.13. A correctable region is also dressed-cleanable.

Proof. Suppose R is correctable. Then for each bare logical operator, such as P_L , there exists a bare CSP operator supported in R^c that implements it. This is also a dressed CSP operator, satisfying the condition for dressed cleanability. ■

We remark again that due to the dependence on dressed CSP operators, Definition II.12 is defined relative to a particular choice of logical subsystem. Similarly to Lemma II.3, a corresponding property holds for dressed-cleanable regions.

Lemma II.14. Suppose a region R is dressed-cleanable. Then any bare CSP operator supported on R implements a logical operation of the form $c1_L \otimes 1_J$, with $c \in \mathbb{C}$.

FIG. 4. Bare and dressed cleanability in tensor network notation. (a) Tensor network depiction of the definition of a bare-cleanable region R . (b) Tensor network depiction of the definition of a dressed-cleanable region R . L denotes the protected logical subsystem, J denotes the junk subsystem, and $\mathcal{P}$ denotes the Pauli group.

Proof. Suppose R is dressed-cleanable. Then take any bare CSP operator A supported on R with $V^\dagger A V = A_L \otimes \mathbb{1}_J$. By assumption, for any $P_L \in \mathcal{P}(\mathcal{H}_L)$, one can find P supported on R^c such that $V^\dagger P V = P_L \otimes P_J$. Since A and P are supported on complementary regions, we have $[A, P] = 0$, which implies $[A_L, P_L] = 0$. Since this is true for any P_L , A_L must commute with all Pauli operators in $\mathcal{P}(\mathcal{H}_L)$; thus, it commutes with all of $\mathcal{B}(\mathcal{H}_L)$, so it must be proportional to the identity, $A_L = c\mathbb{1}_L$. ■

Furthermore, although dressed cleanability is generally a weaker requirement than correctability, these properties are equivalent for a region obeying complementary recovery.

Lemma II.15. *Suppose that some region R obeys complementary recovery, and it is dressed-cleanable with respect to the logical subsystem given by S . Then it is also correctable with respect to S .*

Proof. By assumption, $\mathcal{E}[R^c] \cup \mathcal{E}[R] = \{1, \dots, n_L\}$, the full set of all local encoded subsystems; furthermore, R is correctable with respect to $\mathcal{E}[R^c]$ and R^c is correctable with respect to $\mathcal{E}[R]$. If any local encoded subsystems of S are in $\mathcal{E}[R]$, then correctability of R^c ensures that logical operators on those local encoded subsystems can be implemented on R , which would be nontrivial bare CSP operators, contradicting Lemma II.14. Thus, $S \cap \mathcal{E}[R] = \emptyset$, and therefore $S \subset \mathcal{E}[R^c]$. Since R is correctable with respect to $\mathcal{E}[R^c]$, it is also correctable with respect to S . ■

In stabilizer subsystem codes, it has been shown that the converse directions for Lemmas II.3 and II.14 are also true (see Ref. [54], which is easily generalized from qubit to qudit stabilizer codes). This proof is more involved, however, and fortunately we do not need it here. There is more to be said about a number of different correctability properties of regions of subsystem codes, which we elaborate on further in Appendix A for the interested reader.

E. Fault-tolerant operations in stabilizer subsystem codes

A core result from the theory of quantum error-correcting codes is the Eastin-Knill theorem [27], which states that no code can allow a universal set of gates to be implemented transversally as long as it can detect errors on a single subsystem. Several related results have strengthened this in various contexts [28,29,31,32,55–58], such as more narrowly specifying the set of gates that *can* be implemented transversally. For a family of stabilizer subsystem codes with locally generated stabilizer groups known as *topological stabilizer subsystem codes*, the Pastawski-Yoshida theorem [31] establishes a restriction on transversally implementable gates that depends

only on the dimension of the code's geometry. Specifically, in D dimensions a transversal gate must implement a logical operation to the D th level of the so-called *Clifford hierarchy*, which is defined as follows.

Definition II.16. *Denote $\mathcal{P}$ as the Pauli group on the logical subsystem $\mathcal{H}_L$ and the zeroth level of the Clifford hierarchy to be $\mathcal{C}_0 := \mathbb{C}$. Then the Clifford hierarchy for $i > 0$ is a family of sets $\mathcal{C}_i$ of unitary operators defined recursively as follows [59]:*

$$\mathcal{C}_i = \{U | \forall P \in \mathcal{P}, U P U^\dagger P^\dagger \in \mathcal{C}_{i-1}\}. \quad (6)$$

We emphasize that this is defined in terms of the Pauli strings acting on the logical subsystem $\mathcal{H}_L$.

In any code with a finite distance, another set of fault-tolerant gates is those that do not increase the support of an operator too much under conjugation. For example, suppose that an operator U has the property that when conjugating any local one-site operator X , the resultant operator $U X U^\dagger$ has support on only $k < d$ sites. Then any one-site error will remain correctable after the application of this gate. However, successive applications of these gates may expand the support of the error exponentially, unless there is some particular structure in *how* the errors are spread (or $k = 1$, for example with transversal gates).

Fortunately, when there is a geometric structure associated with physical space, such as in topological codes and holographic codes, we can restrict ourselves to a more specific class of gates with even better protection properties. If application of a unitary U causes the support of an error to spread by some bounded amount, not just in terms of the number of sites affected, but also in terms of the *geometric extent* of these errors, then we say that U is *locality-preserving*. For example, evolution by a local Hamiltonian for a short time can be well approximated by such a unitary (e.g., via Trotterization). Unlike the previous example, successive applications of locality-preserving gates in flat geometries spread errors to only a polynomial number of total sites in the number of applications (raised to the power D , the geometric dimension) rather than an exponential number.

We now formalize this notion.

Definition II.17. *The extent to which a gate U is locality-preserving is quantified by its spread s_U , the maximum distance by which it can increase the support of a local operator supported on some region R . Formally, the spread of U is*

$$s_U = \min\{d \in [0, \infty] | \text{for all regions } R \text{ for all } X \in \text{supp}(R), \\ \times U X U^\dagger \in \text{supp}[B(R, d)]\}. \quad (7)$$

The spread of a transversal operator is zero.

It was shown by Pastawski and Yoshida [31] that for any D -dimensional topological subsystem stabilizer code satisfying some basic error-correction properties [60], any dressed CSP operator with sufficiently limited spread can implement only a dressed logical operator $A_L \otimes A_J$ such that $A_L \in \mathcal{C}_D$. This generalizes a result due to Bravyi and Koenig [28] that applies only to standard topological stabilizer codes. A key step in the proof of this result was the following lemma, which we frequently use in this work.

Lemma II.18 ((Pastawski-Yoshida lemma)). *For a subsystem stabilizer code, let U be a dressed CSP unitary operator supported on the union of regions $\{R_j\}_{j \in [0,m]}$. If R_0 is correctable and each $R_j^+ := B(R_j, 2^{m-j} s_U)$ is dressed-cleanable for $j > 0$, then the logical unitary implemented by U belongs to $\mathcal{C}_m$.*

For completeness, we replicate the proof of this lemma here from Ref. [31].

Proof. Use induction on m . For $m = 0$, U is a dressed CSP operator supported on a correctable region, so by Lemma II.3, it must implement a scalar multiple of the identity on the logical subsystem (i.e., an element of $\mathcal{C}_0$); thus, the base case holds.

Now assume the case m holds and that we wish to show it for $m + 1$. Suppose we have a dressed CSP unitary U supported on $\bigcup_{j=0}^{m+1} R_j$ such that $R_j^+ := B(R_j, 2^{m+1-j} s_U)$ is dressed-cleanable for each $j > 0$. Consider an arbitrary Pauli operator on the logical subsystem, $P_L \in \mathcal{P}$. Because R_1^+ is dressed-cleanable, P can be implemented by a dressed CSP operator P supported in $\text{supp}(P) \subseteq (R_{m+1}^+)^c$. Then

$$\begin{aligned} \text{supp}(UPU^\dagger P^\dagger) &\subseteq \text{supp}(UPU^\dagger) \cap \text{supp}(PU^\dagger P^\dagger), \\ &\subseteq B[\text{supp}(P), s_U] \cap \text{supp}(U), \\ &\subseteq B[(R_{m+1}^+)^c, s_U] \cap \bigcup_{j=0}^{m+1} R_j, \\ &\subseteq R_{m+1}^c \cap \bigcup_{j=0}^{m+1} R_j, \\ &\subseteq \bigcup_{j=0}^m R_j, \end{aligned} \quad (8)$$

where we used the fact that $R_{m+1}^+ = B(R_{m+1}, s_U)$ in the second-last line. We can now apply the inductive hypothesis to the dressed CSP unitary $UPU^\dagger P^\dagger$ so long as $B(R_j, 2^{m-j} s_{UPU^\dagger P^\dagger})$ are dressed-cleanable. Fortunately, since $s_{UPU^\dagger P^\dagger} \leq 2s_U$, this is implied by the fact that $B(R_j, 2^{m+1-j} s_U)$ is dressed-cleanable, and thus $UPU^\dagger P^\dagger \in \mathcal{C}_m$. By the definition of the Clifford hierarchy, we have $U \in \mathcal{C}_{m+1}$. ■

This lemma is simpler in the case that U is transversal, which implies that $s_U = 0$ and so $R_j^+ = R_j$. We also introduce a stronger form of this lemma in the case that U is

also bare CSP. We prove it using a proof virtually identical to that of Lemma II.18 from Ref. [31].

Lemma II.19. *For a subsystem stabilizer code, let U be a bare CSP unitary operator supported on the union of regions $\{R_j\}_{j \in [0,m]}$. If R_0 and each $R_j^+ := B(R_j, 2^{j-1} s_U)$ are dressed-cleanable, then the logical unitary implemented by U belongs to $\mathcal{C}_m$.*

Proof. Use induction on m . For $m = 0$, U is a bare CSP operator supported on a dressed-cleanable region, so by Lemma II.14, it must implement a scalar multiple of the identity on the logical subsystem (i.e., an element of $\mathcal{C}_0$); thus, the base case holds.

The proof then follows identically to that of Lemma II.18, until it comes time to apply the inductive hypothesis, which now additionally requires that $UPU^\dagger P^\dagger$ be a bare CSP operator. Fortunately, as U is a bare CSP operator, even though P is only a dressed CSP operator, the component P_J acting on the junk subsystem cancels with its inverse, and the product is indeed a bare CSP operator. ■

The final lemma of this section establishes a useful sufficient condition, in terms of price, to guarantee the existence of three correctable regions that suffice to place restrictions on the logical action of fault-tolerant gates.

Lemma II.20. *If a logical subsystem S of a subsystem stabilizer code satisfies the relationship*

$$2 \leq p(S) \leq 2d(S) - 2, \quad (9)$$

then any transversal dressed CSP unitary U can implement only an element of the Clifford group $\mathcal{C}_2$ on the logical subsystem associated with S .

Proof. By the definition of price, there exists some region R with size $|R| = p(S)$ such that R^c is correctable. Then as $|R| \geq 2$, R can be split into two nonempty regions R_1 and R_2 such that $|R_1|, |R_2| \leq d(S) - 1$, meaning that R_1 and R_2 are smaller than the smallest noncorrectable region; thus, they must be correctable. Also, recall that Lemma II.13 showed that correctability is stronger than dressed cleanability, so these regions are also dressed-cleanable. Thus, we can apply Lemma II.18 with $s_U = 0$ because of its transversality to conclude that U implements a Clifford gate, $U_L \in \mathcal{C}_2$. ■

We re-express the lemma in this way to crystallize the physical intuition for why holographic codes should not permit transversal non-Clifford gates for sufficiently small bulk regions. In a code satisfying complementary recovery, it was shown in Ref. [39] that the distance is equal to the price for a logical subsystem associated with a single point. Although holographic stabilizer codes do not generally satisfy complementary recovery, we use this observation as

motivation to argue that the difference between price and distance is an effective measure of the effective “size” of a region, after accounting for the degree to which complementary recovery fails. In this way, Lemma II.20 can be interpreted as saying non-Clifford gates cannot be transversally implemented on sufficiently small regions, i.e., those with $p - d \leq d - 2$.

An analogous definition can be formed in terms of the generalized notions of price and distance, $p_\sigma(S)$ and $d_\sigma(S)$. Finally, the lemma could also be strengthened by redefining distance in terms of the smallest *non-dressed-cleanable* region; this would be a larger notion of distance and thus easier to satisfy Eq. (9). R_1 and R_2 would be dressed-cleanable, and thus Lemma II.18 could still be invoked.

Lemmas II.18 and II.19 were both originally proven in the context of qubit stabilizer codes, but easily generalize to qudits. From here on, we restrict our exposition to the case of qubits for simplicity; however, we emphasize that all of our results apply likewise to qudit-based holographic stabilizer codes.

III. HaPPY CODES AND NON-CLIFFORD GATES

The simplest holographic stabilizer code, commonly referred to as “HaPPY” [6], is built out of *perfect tensors*—a class of tensors with particularly nice error-correction properties that allow the entanglement wedge map to mimic the geometric behavior of AdS/CFT. As the encoding map is given by a tensor network built from identical local encoding tensors, any gates that can be implemented transversally on this encoding tensor are inherited by the larger network [61]. For example, the standard choice of perfect tensor is the encoding isometry of the five-qubit code (see the appendix in Ref. [6] for a review). This code admits a transversal gate set generated by the homogeneous tensor product of single-site Pauli group elements and the Clifford gate K , which can be defined as mapping $X \rightarrow Z \rightarrow Y \rightarrow X$ under conjugation [62]. Transversal implementation of the group generated by these gates is inherited by the larger network built out of this perfect tensor, as shown in Fig. 5. This is the most straightforward way in which a holographic code can allow transversal implementation of a logical gate; and we now argue that no non-Clifford gate can be implemented in such a fashion.

Consider a rank- $2N$ tensor $T_{i_1, i_2, \dots, i_{2N}}$, and suppose that for any $n \leq N$, the matrix given by lowering the first n indices is an isometry, i.e.,

$$T_{i_1, \dots, i_n}^{j_{n+1}, \dots, j_{2N}} T_{i_{n+1}, \dots, i_{2N}}^{j_1, \dots, j_n} = \delta_{i_1, \dots, i_n}^{j_1, \dots, j_n}. \quad (10)$$

If this holds not only for the first n indices but for *any* choice of n indices, then we say that T is a perfect tensor. The tensor need not be permutation invariant, i.e., it

FIG. 5. Pictorial argument that transversal implementation of a gate on a single perfect tensor implies transversal implementation for the entire HaPPY network.

may form different isometries depending on which legs are chosen to be inputs and which are chosen to be outputs. For a review of perfect tensors, see Ref. [6].

Such a tensor can be interpreted as a good quantum error-correcting code with parameters $[[n, k, d]] = [[2N - 1, 1, N]]$ by lowering only the first index and treating this as the encoding isometry:

$$V := T_{i_1}^{i_2, \dots, i_{2N}} : \mathcal{H}_L = \mathcal{H}_2 \rightarrow \mathcal{H}_P = \mathcal{H}_2^{\otimes (2N-1)}, \quad (11)$$

with the distance of N being guaranteed by the perfect tensor property (see Ref. [6] for details). Furthermore, it guarantees that any region of $N - 1$ qubits or fewer is correctable. Thus, the $2N - 1$ qubits in the physical system can easily be partitioned into three sets such that each is correctable. By Lemma II.18, this means that for any transversal unitary U , the logically implemented gate is a Clifford gate, $V^\dagger U V \in \mathcal{C}_2$.

Furthermore, the same can be said about holographic codes built from a slightly broader class of tensors known as *planar perfect* [63] or *block perfect* [11] tensors. In these tensors, the order of the indices matters—the isometry property is required to hold only for any set of n consecutive indices (where i_{2N} is understood to be adjacent to i_1).

Of course, the above argument does not rule out the possibility of transversally implemented non-Clifford gates in networks built out of these tensors—it excludes only the simplest possible way for such a gate to be implemented, which would be that it is transversal on each individual tensor that makes up the network. Nevertheless, rigorous arguments can be formulated that rule out any kind of transversal non-Clifford in a HaPPY code. Essentially, one starts with a logical operator on some central tensor

and considers the boundary regions that it can be pushed out to at each successive layer of the network. Because of the “negative curvature” (which is a necessary requirement for the network to form a bulk to boundary isometry), the angular size of the region to which it can be pushed does not increase too much layer by layer; the total size of each layer grows at least as quickly as the region it is being pushed to does. Thus, the complement of the region to which it is being pushed, which is correctable, is sufficiently large that one can divide the physical space into three such regions and proceed by applying Lemma II.18 to exclude the possibility of transversal non-Clifford gates.

In Appendix B, we use a rigorous form of the above argument to prove that any HaPPY network built out of a homogeneous hyperbolic planar tiling does not accommodate the transversal application of non-Clifford gates. The idea behind our proof of no transversal non-Clifford gates in holographic tensor network stabilizer codes can be extended to rule out more general locality-preserving gates. In the following section we describe such a generalization for more general families of sufficiently holographic stabilizer codes that should include tensor network codes as a special case.

IV. COMPLEMENTARY RECOVERY AND NON-CLIFFORD GATES

In the previous section we argued that HaPPY codes generally cannot admit transversal non-Clifford gates. In this section we turn our attention to wider families of holographic codes that have been introduced and studied more recently. The natural question that occurs is whether the transversal Clifford restriction is a particular attribute of HaPPY codes, or a more general phenomenon due to some underlying physical reason that extends to these other families. In Sec. IV A, we argue for the latter—specifically that it is a consequence of any code for which regions satisfying complementary recovery are sufficiently common. The argument we present can even apply if no regions satisfy complementary recovery, as long as some almost do (see Sec. IV C). We extend our consideration to finite-size regions of the bulk in Sec. IV B, codes with only approximate stabilizer encoding maps in Sec. IV D, nontransversal locality-preserving gates in Sec. IV E, and codes with nontrivial centers in Sec. IV F.

A. Complementary recovery and transversal non-Clifford gates

First, we begin with a proof that satisfying complementary recovery for every region is sufficient to exclude transversal non-Clifford gates. A stronger variant of this theorem is one in which only every *connected* region needs to satisfy complementary recovery. This is stronger for two reasons: firstly, the assumption is less stringent and applies

to more codes, and secondly, the assumption of $d_c \geq 2$ is weaker than that of $d \geq 2$.

Theorem IV.1. *Consider a holographic stabilizer code with encoding isometry V and a particular choice of local bulk subsystem i . Suppose that every (connected) region satisfies complementary recovery, and the (connected) code distance is at least $d_{(c)}(\{i\}) \geq 2$. Then any transversal dressed CSP unitary U can act on this subsystem only via an element of the Clifford group, $VUV^\dagger = U_L \otimes U_J$, $U_L \in \mathcal{C}_2$.*

Proof. Let the region R be the smallest noncorrectable region. As R is noncorrectable, the entanglement wedge of its complement $\mathcal{E}[R^c]$ cannot contain i . By complementary recovery, this means that $i \in \mathcal{E}[R]$, and so R^c must be correctable. Since we have a (connected) region of size $d_{(c)}(\{i\})$ with a correctable complement, the price must be at most $d_{(c)}(\{i\})$, and therefore $p_{(c)}(\{i\}) = d_{(c)}(\{i\})$ using Lemma II.10. Thus, by assumption, $2 \leq p_{(c)}(\{i\}) \leq 2d_{(c)}(\{i\}) - 2$. By Lemma II.20, any transversal dressed CSP unitary can implement only an element of the Clifford group, $\mathcal{C}_2$. ■

Of course, insisting on complementary recovery for every region, or even every connected region, is highly restrictive. We do not even know if such a holographic stabilizer code exists [64]. Fortunately, the following theorem requires a much weaker sufficient condition to exclude the possibility of transversal implementation of non-Clifford gates on a logical subsystem. The idea is that we need just one region R satisfying complementary recovery to exist that would be suitable for use in the proof of Theorem IV.1. This region R need satisfy only a few basic properties for the proof to work.

Theorem IV.2. *Consider a holographic stabilizer code with encoding isometry V and a particular choice of local bulk subsystem i . Suppose there exists a noncorrectable (connected) region R with size $2 \leq |R| \leq 2d_{(c)}(\{i\}) - 2$ that obeys complementary recovery [65]. Then any transversal dressed CSP unitary U can act on this subsystem only via an element of the Clifford group, $VUV^\dagger = U_L \otimes U_J$, $U_L \in \mathcal{C}_2$.*

Proof. The proof proceeds exactly as that of Theorem IV.1. ■

An example of the construction used in this proof is shown in Fig. 6. In many holographic codes, such as certain instances of HaPPY, we expect the connected code distance to be significantly larger than the code distance, making the connected form of this theorem again more useful [66].

FIG. 6. An alternating pentagon-hexagon HaPPY code. We show how to construct three appropriate regions to be used in the proof of Theorem IV.2. In this case, the logical subsystem consists of only the central bulk qubit, and its connected code distance d_c is 13; an example of a connected noncorrectable region R with this size, along with its entanglement wedge, is highlighted in yellow. In this case, R obeys complementary recovery, and hence can be used to prove Theorem IV.2 as follows: because of complementary recovery, R^c (highlighted along with its entanglement wedge in blue), is correctable. Furthermore, the two regions highlighted in red must be correctable as they are smaller than R , the smallest noncorrectable region; thus, the boundary is partitioned into three correctable regions. Lemma II.18 then implies that a non-Clifford gate cannot be applied transversally on the central qubit.

In any case, the definition of (connected) code distance implies the existence of at least one noncorrectable region of size $d_{(c)}$, in addition to a number of (connected) supersets of this region that are also noncorrectable and smaller than $2d_{(c)} - 2$. As long as regions obeying complementary recovery are not too rare, which they should not be for a purported holographic code, at least one such region will obey complementary recovery and the theorem will apply. For example, although HaPPY codes do not generally obey complementary recovery, Fig. 6 shows an example of a HaPPY code in which an appropriate region can still be found to apply Theorem IV.2.

One can also construct a variant for any definition of size $\sigma(R)$, as per Definition II.8, so long as it is always possible to break a region into two smaller regions according to this notion of size. As long as the assumptions are satisfied for any such notion of size, the conclusion holds.

One can strengthen this theorem by noting that the three regions do not necessarily need to be correctable. Two regions need only be dressed-cleanable, which is a weaker condition in general—relaxing the assumptions necessary to apply the theorem. For example, the red and blue regions of the HaPPY code in Fig. 7 are both dressed-cleanable but not correctable. This freedom is irrelevant in a code

FIG. 7. In this case, with a two-site logical region, we can make an even stronger claim about what can be implemented by a bare CSP operator using Lemma II.19. The arrows denote the directions in which logical operators are “pushed” using the perfect tensor property; vertical arrows on a bulk tensor show that the operator is pushed to have support on that bulk leg, which acts as a dressing. Thus, this pushing protocol shows that the complements of both the red solid region and the blue solid region are dressed-cleanable, as any logical operator can be supported on the boundary up to some dressing on the junk subsystem. Thus, they can be used in Lemma II.19 with $m = 1$, meaning that a bare CSP operator can apply a Pauli operator only to the central qubits; even a non-Pauli Clifford gate is forbidden.

with all regions satisfying complementary recovery, due to Lemma II.15—which also suggests holographic codes should not generally have too large a difference between dressed-cleanable and correctable regions.

It follows from the proof that if R is close to $d_{(c)}$ in size, then there is significant “room to spare” when dividing the boundary system into three correctable regions. That is, R_1 and R_2 are not only smaller than the smallest noncorrectable region, they are roughly half its size. Instead of choosing R_0 , R_1 , and R_2 to be mutually nonintersecting, we could instead choose them to have some overlap while still being correctable and satisfying $R = R_0 \cup R_1 \cup R_2$. Such an overlap proves useful in many of our subsequent extensions of this simple initial setup.

For example, it proves useful in generalizing beyond the case where the logical subsystem $S \subset \{1, \dots, n_L\}$ is chosen to consist of only a single local encoded subsystem i . First, we point out why the argument as presented above breaks down if multiple local encoded subsystems are included. For example, suppose that two local encoded subsystems are contained in the logical subsystem, $S = \{i, j\}$. Then it may be the case that $i \in \mathcal{E}[R^c]$ and $j \in \mathcal{E}[R]$, in which case both R and R^c are noncorrectable with respect to $\{i, j\}$ even if complementary recovery is satisfied. By contrast, with a single local encoded subsystem, complementary recovery guarantees that either R or R^c is correctable. Nonetheless,

FIG. 8. An example with a logical region containing more than one site (each labeled L), illustrating Theorem IV.3. More perfect tensors are included than usual so as to be able to pick boundary regions with greater precision. A noncorrectable boundary region R has an entanglement wedge (enclosed by a solid black line) that borders the logical region. R needs to be expanded only slightly, to R_d , to include the whole of the logical region in its new entanglement wedge (enclosed by a dashed black line). Thus, R_d^c is correctable. Furthermore, we can break R_d into two correctable regions, which are highlighted, along with their in entanglement wedges, in yellow and red. Thus, we find three correctable regions and Lemma II.18 can be applied.

in many cases, three appropriate regions can still be constructed for a collection of sites in the bulk; an example in a HaPPY code is shown in Fig. 8. We argue below that due to the existence of room to spare in the above argument, this should generally be the case.

B. Finite-size bulk regions

We now generalize our argument to bulk regions with a finite, nonzero size. First we attempt to provide some intuition. We have seen that a logical subsystem associated with a point in a holographic code satisfies (connected) distance equals (connected) price. For extended bulk regions this equality no longer holds. However, from AdS/CFT we expect for the logical subsystem associated with a ball of radius δ in the bulk $p_c \leq d_c + \kappa(\delta)$, where $\kappa(\delta)$ grows exponentially with δ for any spatial dimension. Hence, for sufficiently small bulk regions of a code with a holographic structure, we expect $2 \leq p_c \leq 2d_c - 2$ is still satisfied. By Lemma II.20 this implies that transversal gates can act only via Clifford gates on the logical subsystem associated with any such bulk region. This argument is extended to locality-preserving gates for logical regions that are sufficiently deep within the bulk below.

To make this generalization more rigorous, we need to refer to the geometric structure of the code, as introduced in Sec. II C. Specifically, we use the notion of a

κ_R -smooth entanglement wedge map. As discussed where it was introduced, κ_R is typically exponentially large in r ; however, the amount of “room to spare” available from the proof of Theorem IV.2 is proportional to the code distance d , which is also expected to be exponential in r for deep bulk inputs. Thus, we are usually able to trade this overlap for the inclusion of a size $\mathcal{O}(1)$ region in the bulk.

Theorem IV.3. *Consider a holographic stabilizer code with encoding isometry V and a particular choice of local bulk subsystem S of diameter D_S with a κ_R -smooth entanglement wedge. Suppose there exists a noncorrectable connected region R with size $2 \leq |R| \leq 2d_c(S) - 2 - \kappa_R(D_S)$ that obeys complementary recovery. Then any transversal dressed CSP unitary U is constrained to act on this subsystem via an element of the Clifford group, $VUV^\dagger = U_L \otimes U_J$, $U_L \in \mathcal{C}_2$.*

Proof. As R is not correctable and obeys complementary recovery, there is at least one subsystem $i \subset S$ contained in $\mathcal{E}(R)$. $S \subset B(\mathcal{E}(R), D_S)$ because the distance between two points in S is at most D_S . Thus, if $x = \kappa_R(D_S)$, then there exists a connected boundary region R_x such that $S \subset \mathcal{E}(R_x)$, meaning its complement R_x^c is correctable. Thus, the code price is upper bounded by $|R_x| = |R| + x$, which in turn is upper bounded by $2d_c(S) - 2$, so we can apply Lemma II.20 to conclude the proof. ■

C. Almost complementary recovery

We now move on to consider a more general case where the entanglement wedge map is smooth but the region R only almost obeys complementary recovery. In this case, i may not belong to the entanglement wedge $\mathcal{E}[R]$, but a slightly larger $R^+ \supset R$ can be chosen such that $i \in \mathcal{E}[R^+]$, and thus R^{+c} is correctable with respect to $\{i\}$.

Definition IV.4. *We say that a region R satisfies ε -almost complementary recovery if it can be expanded by ε in all directions, such that the new entanglement wedge includes the complement of the entanglement wedge of R^c , i.e., $\mathcal{E}[R^c]^c \subset \mathcal{E}[B(R, \varepsilon)]$ [67].*

Theorem IV.5. *Consider a holographic stabilizer code with encoding isometry V and a particular choice of local bulk subsystem S of diameter D_S with a κ_R -smooth entanglement wedge. Suppose there exists a noncorrectable connected region R with size $2 \leq |R| \leq 2d(S) - 2 - \kappa_R(D_S) - \varepsilon$ that obeys ε -complementary recovery. Then any transversal dressed CSP unitary U can act on this subsystem only via an element of the Clifford group, $VUV^\dagger = U_L \otimes U_J$, $U_L \in \mathcal{C}_2$.*

Proof. The proof is a straightforward combination of Definition IV.4 and the proof of Theorem IV.3. ■

D. Approximate stabilizer codes

AdS/CFT is believed to not be an exact quantum error-correcting code, as its error-correction properties hold only to first order in a perturbative expansion [5]. Instead, it is best understood as an approximate code [5]. One family of toy models [9] attempts to capture this aspect of AdS/CFT in terms of an encoding map that is only *approximately* equal to an encoding isometry of a stabilizer code. Specifically, the map $V: \mathcal{H}_L \rightarrow \mathcal{H}_P$ is given by

$$V = V_0 + \delta V, \quad (12)$$

with V_0 an encoding isometry for a stabilizer code and δV some small perturbation; $\|\delta V\|_1 = \varepsilon$ for some small $\varepsilon > 0$.

In such a case, it is straightforward to see that a CSP unitary must implement a logical gate which is approximately equal to an element of the Clifford group. To see this, note that V_0 is an exact holographic stabilizer code, and thus the unitary applied to this code implements some element of the Clifford group, $V_0^\dagger U V_0 = U_L \in \mathcal{C}_2$. Then one can show using Hölder's inequality that

$$\begin{aligned} \|V^\dagger U V - U_L\|_1 &= \|(V_0 + \delta V)^\dagger U (V_0 + \delta V) - V_0^\dagger U V_0\|_1, \\ &\leq \|V_0 U \delta V^\dagger\|_1 + \|\delta V U V_0^\dagger\|_1 + \|\delta V^\dagger U \delta V\|_1, \\ &\leq \|\delta V\|_1 + \|\delta V\|_1 + \|\delta V\|_1 \|\delta V^\dagger\|_\infty, \\ &\leq 2\varepsilon + \varepsilon^2, \end{aligned} \quad (13)$$

i.e., the logical operation implemented by the full code V is close to an element of the Clifford group.

E. Nontransversal locality-preserving gates

We now combine the Pastawski-Yoshida lemma, Lemma II.18, with the geometric structure introduced in Sec. IIC to extend the previous result to the case of unitaries that are not necessarily transversal but are still locality-preserving (i.e., that have some finite spread $s_U > 0$).

This generalization is simplest to formulate when the spatial boundary geometry is one dimensional.

Theorem IV.6. *For a subsystem stabilizer code S with a one-dimensional boundary, if there is a noncorrectable region R of size $d_c(\{i\})$ that satisfies complementary recovery, then any dressed CSP unitary operator U with $s_U < (d_c/6)$ implements a logical unitary in $\mathcal{C}_2$ on the logical system $\{i\}$.*

Proof. To apply Lemma II.18, we need to partition the boundary space into three regions such that one is correctable, one is correctable even after expansion by a distance of s_U on either side, and one is correctable after expansion by $2s_U$ on either side. Furthermore, the

definition of d_c means that any connected region smaller than d_c is correctable. Thus, we can proceed as follows.

Parametrize the interval R as $[0, d_c]$ and choose R_1 as $[0, (2d_c/3)]$ and R_2 as $[(2d_c/3), d_c]$. Then for any unitary with spread $s_U < (d_c/6)$, the extended regions $R_1^+ = B(R_1, s_U) = [-s_U, (2d_c/3) + s_U]$ and $R_2^+ = B(R_2, 2s_U) = [(2d_c/3) - 2s_U, d_c + 2s_U]$ are both less than d_c in length, and therefore remain correctable; R_0 also remains correctable, and thus Lemma II.18 applies. ■

More generally, this argument extends to higher dimensions as follows: Let d be the size of the smallest non-correctable region. Let R be a connected region of size d . We then split R into two regions R_1 and R_2 of sizes $(2d/3)$ and $(d/3)$, respectively. We assume that we can choose sufficiently nice regions R_1 and R_2 in the sense that the sizes of their extensions $B(R_i, s)$ do not increase too fast [68] with s . Then for sufficiently small s_U , the regions $R_i^+ = B(R_i, 2^{2-i}s_U)$ have size smaller than d and thus remain correctable. We can then apply Lemma II.18 to achieve the desired result. Following this line of reasoning, Theorem IV.6 generalizes straightforwardly to a higher-dimensional setting where the boundary is a sphere and the region R is a ball. At this time we do not know of a clean statement of the most general higher-dimensional version of our result. This is in part due to the richer set of possibilities for the topology of connected regions in dimensions higher than 1. Furthermore, degenerate cases exist where the region R is itself a (thickened) boundary and therefore the scaling of $B(R_i, s)$ with s for any subregion is necessarily too fast to accommodate any positive spreading $s_U > 0$. However, recall that the region R was chosen as a noncorrectable region with minimal size. In AdS/CFT, such regions are generally expected to be solid spheres [69], and therefore the analogous regions in discrete models of AdS/CFT are expected to be approximately spherical and thus avoid the troublesome degenerate cases mentioned above.

F. Entanglement wedge surface algebras with nontrivial centers

Subsystem codes are a special case of *operator-algebraic quantum error correction* in which some general algebra of operators (the “logical algebra”) is protected against error. In subsystem codes, this is the algebra of bare logical operators $\{A_L \otimes \mathbb{1}_J\}$. An algebra of this form is special in that its only central elements are those proportional to the identity. More generally, the logical algebra may have a *nontrivial center* (i.e., it may contain elements that commute with all other elements but are not proportional to the identity on L). The physical consequence of this is that the logical region is no longer associated with a Hilbert space—it is now identified with a logical *algebra* instead [70,71]. Because the entanglement wedge map associates a bulk region with each boundary region, this means it

now identifies each boundary region with an algebra and not a subsystem [5]. As a result, the limit from Eq. (2) no longer applies, and the regions $\mathcal{E}[R]$ and $\mathcal{E}[R^c]$ can have a nonempty overlap $\mathcal{E}(R) \cap \mathcal{E}(R^c) \neq \emptyset$, which we refer to as the “entangling surface.” Because this algebra can be reconstructed on independent regions of the boundary, it must be associated with an Abelian algebra, which also forms the center of the algebras associated with $\mathcal{E}[R]$ and $\mathcal{E}[R^c]$, respectively.

For example, in Ref. [8] codes were introduced (which we refer to as “LOTE codes” after the title of that paper) that are constructed from HaPPY codes but with additional qubits associated with the entangling surface of any given region R . However, only the algebra of operators generated by $\{1, \sigma_Z\}$ acting on such a qubit can be recovered on either R or R^c . The full set of Pauli operators $\{1, \sigma_X, \sigma_Y, \sigma_Z\}$ can still be recovered on the boundary, but not on R or R^c alone. This means that with respect to this partition of the boundary, the data on the entangling surface are composed of classical bits. This is meant to model an intuitive feature of AdS/CFT: the area of the entangling surface is information that ought to be recoverable from both R and R^c .

Even for LOTE codes we are still able to apply our methods to exclude transversal implementation of non-Clifford gates. We do this fairly straightforwardly—we simply ignore the additional algebraic structure, and proceed as if it were a subsystem code with local encoded subsystems as in Eq. (1). A qubit on the entangling surface has a subalgebra reconstructable on either R or R^c , but its entire algebra can be reconstructed on neither, and thus it does not belong to the subsystem associated with either region’s maximal entanglement wedge. In other words, such qubits contribute to any given region’s failure to obey complementary recovery. Thus, there are no nontrivial regions satisfying complementary recovery (as we have defined it here) in such a code. Nonetheless, so long as regions satisfying complementary recovery are not too rare in the underlying HaPPY code, regions that *almost* satisfy complementary recovery are not too rare in the corresponding LOTE code. Thus, we can apply the results of Sec. IV C to show that non-Clifford gates cannot be implemented in the *interior* of a region’s entanglement wedge. It is unclear how to define a Clifford operator on an algebra with a mixture of bits and qubits, so this appears to be the most that can be said—locality-preserving gates cannot implement logical non-Clifford gates on a slightly smaller quantum subalgebra that excludes the classical parts.

G. Other levels of the Clifford hierarchy

For the special case that U is a bare CSP operator, Lemma II.19 implies a stronger result; namely, that one can allow R_0 to only be dressed-cleanable rather than correctable. This can be used to strengthen the result further. In some cases it can even be used to show that only Pauli

gates can be implemented, such as in Fig. 7, by dividing the boundary into just two dressed-cleanable regions R_0 and R_1 . This case of the bound is relevant if one wishes to implement a desired bare logical gate on a small bulk subsystem without causing any disturbances to other regions of the bulk.

It is not possible to restrict a dressed CSP operator to implement something only from $\mathcal{C}_1$ in this way. To do so would require expressing the boundary as a union of a dressed-cleanable region and a correctable region, which is impossible; any bare logical operator could be cleaned to R_1 using correctability of R_0 , but Lemma II.14 applied to R_1 means that the operator must be trivial—a contradiction.

Finally, for sufficiently large bulk regions, the arguments of Sec. IV E do not apply. The reason for this is that all correctable regions are too small to cover the entire boundary using only three of them; a greater number would be required. In these cases, a restriction on locality-preserving gates still exists, depending on the number k of required correctable regions. Thus, even if a non-Clifford gate cannot be ruled out for very large bulk regions, some restriction still exists corresponding to a higher level of the Clifford hierarchy, $\mathcal{C}_{k-1}$.

Theorem IV.7. *For a subsystem stabilizer code S with a one-dimensional boundary, if there is a region R of size $d_c(S)$ that satisfies complementary recovery, then any dressed CSP unitary operator U with $s_U < \{d_c/[2(2^m - 1)]\}$ implements a logical unitary in $\mathcal{C}_m$ for $m \geq 2$.*

Proof. As in Sec. IV E, we subdivide the region R into m consecutive regions R_j of sizes $(2^{m-j} d_c)/(2^m - 1)$. Then $R_j^+ = B(R_j, 2^j s_U)$ has size $(2^{m-j} d_c)/(2^m - 1) + 2^{j+1} s_U < ([2^{m-j} + 2^{j+1}]/[2^m - 1]) d_c \leq d_c$. We can thus apply Lemma II.18 and obtain the result. ■

The above argument can be generalized to higher dimensions along similar lines to the discussion at the end of Sec. IV E. Again we assume that there is a boundary region R of the smallest noncorrectable region size that can be decomposed into a set of subregions R_i , $i = 1, \dots, m$, that are sufficiently nice in the sense that the sizes of their extensions $B(R_i, s)$ do not grow too fast with s . Then for sufficiently small s_U , the argument presented above for Theorem IV.7 can be applied straightforwardly, leading to the same conclusion. Similarly to Sec. IV E, we do not currently have a clean statement of the most general version of this result but we expect the statement for a solid spherical boundary region to be the most relevant.

V. DISCUSSION AND CONCLUSIONS

In this work we have established a specialization and strengthening of the Eastin-Knill theorem for holographic subsystem stabilizer codes, inspired by analogous work

due to Bravyi and Koenig and Pastawski and Yoshida for topological stabilizer codes. Specifically, we have shown that sufficiently locality-preserving operations on the physical qubits of a holographic stabilizer code can implement only logical Clifford gates on sufficiently small regions of the bulk. This result was extended in several ways, including to approximate stabilizer encodings and to codes in which algebras with nontrivial centers live on the surfaces of entanglement wedges. We further showed that on weakening our assumptions, locality-preserving gates are still restricted to implement an element from a fixed level of the Clifford hierarchy. In the course of deriving our results, we also introduced a general definition of the maximal entanglement wedge for arbitrary subsystem codes which may be of independent interest. While we have focused on the simple case of the hyperbolic disk in our figures, our results also apply to nontrivial bulk topologies and higher-dimensional spaces.

There is a recurring theme in fault-tolerant quantum computation that the set of fault-tolerantly implementable gates is more severely restricted the better the error-correction properties of a code are. This makes some intuitive sense—better protected logical information should also be harder to manipulate. For example, the Eastin-Knill theorem establishing the impossibility of universal transversal gate sets applies only when the code distance is greater than 1; the Bravyi-Koenig theorem and Pastawski and Yoshida’s subsystem generalization require a macroscopically large code distance; and another result obtained by Pastawski and Yoshida restricts transversally implementable gates more strongly the higher the loss threshold is. Similarly, our results require assumptions about the quality of the code; namely, that the distance of a region is “large enough” with respect to the price of the region.

An interesting future direction is the extension of our results restricting locality-preserving gates to the setting of the full AdS/CFT duality, possibly by exploiting conformal invariance. This approach is inspired by the generalization of the Bravyi-Koenig bound to all $(2+1)$ -dimensional topological quantum field theories (on closed manifolds) [29] that was established by exploiting consequences of topological invariance. As explained in Sec. I, it has been successfully argued that AdS/CFT supports no global symmetries [36]. While this result partially overlaps with our goal, our aim is somewhat stronger, restricting the possible bulk evolutions generated by arbitrary locality-preserving and codespace-preserving evolutions on the boundary with no requirement that they take the form of a global symmetry. This could have potentially interesting implications for the prospect of implementing bulk Hamiltonian evolution via locality-preserving operations on the boundary.

There is an apparent tension between the results of this paper and the intuition that there ought to exist states in AdS/CFT that contain a quantum computer in the bulk. Holographic codes built from tensor networks are believed

to capture the error-correcting behavior of AdS/CFT in a semiclassical regime with a fixed geometry (i.e., without gravity being fully quantum), but systems existing in this geometry can still be fully quantum, and thus a quantum computer should be possible. In a short time, this quantum computer ought to be able to perform, at least approximately, an arbitrary element of $SU(2^k)$ on k qubits for $k \geq 2$ (but not much greater). But short-time evolution is implemented on the boundary by a locality-preserving unitary, which we have seen can only implement Clifford gates in holographic stabilizer codes. Of course, in AdS/CFT the local bulk or boundary degrees of freedom are larger than simply qubits, so a fairer comparison should use qudit rather than qubit stabilizer codes. However, even for a large qudit dimension it is impossible for the Clifford group to contain a subgroup arbitrarily approximating $SU(2^k)$ [72]. Thus, it appears unlikely that exact or approximate (as defined in Sec. IV D) stabilizer codes alone can capture this aspect of AdS/CFT. Our results can hence be interpreted as placing severe limitations on how well stabilizer toy models are able to capture aspects of holography to do with locality-preserving boundary evolutions. In future work we plan to show that by adding additional constraints to the logical bulk of a qudit stabilizer code it is indeed possible to build an approximate holographic code that can transversally implement any element of $SU(2^k)$.

It would be very interesting to combine our restrictions on locality-preserving operations based on holographic structure and locality with the approximate covariant code trade-off bounds derived in Ref. [73]. We expect that both of these directions are relevant to studying the action of locality-preserving evolutions in the full AdS/CFT duality, which is expected to lead to approximate quantum codes with a holographic locality structure.

ACKNOWLEDGMENTS

We thank Patrick Hayden for helpful discussions and comments on the manuscript. We also thank Jon Sorce for pointing us to Ref. [69]. K.D. is supported by the Center for Science of Information, an NSF Science and Technology Center, under Grant Agreement No. CCF-0939370. S.C. acknowledges support from the Knight-Hennessy Scholars program. D.W. acknowledges support from the Simons Foundation.

APPENDIX A: CORRECTABILITY PROPERTIES FOR REGIONS OF SUBSYSTEM CODES

In Sec. II, we introduced two correctability properties that a region of a subsystem code can have—correctability and dressed cleanability. In this appendix, we explore a broader family of such properties, and present the known relationships between them.

FIG. 9. Relationships between properties of physical regions in subsystem codes. The blue relationship holds for regions with complementary recovery, such as those in holographic codes.

1. Subsystem codes

To begin with, in general subsystem codes, a notion of dressed-cleanable can be defined analogously to Definition II.12 but for any dressed operator, not just Pauli operators.

Definition A.1. A region R of a subsystem code is defined to be

- (a) *bare-cleanable* (aka “correctable”) if and only if for any logical operator A_L there exists a bare CSP operator A supported on R^c that implements $A_L \otimes \mathbb{1}_J$;
- (b) *dressed-cleanable* if and only if for any logical operator A_L there exists a dressed CSP operator A supported on R^c that implements $A_L \otimes A_J$;
- (c) *bare-trivial* if and only if any bare CSP operator supported on R implements a logical operation of the form $c\mathbb{1}_L \otimes \mathbb{1}_J$, with $c \in \mathbb{C}$;
- (d) *dressed-trivial* if and only if any dressed CSP operator supported on R implements a logical operation of the form $\mathbb{1}_L \otimes A_J$ for some $A_J \in \mathcal{B}(\mathcal{H}_J)$.

Lemmas II.3 and II.14 can then be expressed as meaning that bare-cleanable implies dressed-trivial and dressed-cleanable implies bare-trivial. Furthermore, Lemma II.13 means that bare-cleanable implies dressed-cleanable, and it is easy to show similarly that dressed-trivial implies bare-trivial. We summarize these results, along with Lemma II.15 for regions satisfying complementary recovery, in Fig. 9.

2. Stabilizer subsystem codes

In the special case of a stabilizer subsystem code, there is some additional structure. The definition in Definition A.1 is still valid, but we could instead choose to define analogous properties that require the respective properties to hold only for Pauli operators. It turns out that most of these properties are equivalent to those in Definition A.1, for example:

Lemma A.2. A region R of a stabilizer subsystem code is bare-cleanable if and only if for any logical Pauli operator P_L there exists a bare CSP Pauli operator P supported on R^c that implements $P_L \otimes \mathbb{1}_J$.

Proof. ($\implies$) For any logical operator P_L , bare cleanability tells us there is a bare CSP operator A supported on R^c that implements $P_L \otimes \mathbb{1}$. A can be decomposed into Pauli operators supported on R^c , as $A = \sum_i \lambda_i P^i$. Since A commutes with every stabilizer, each of these Pauli operators P^i must also commute with all the stabilizers, and therefore be codespace-preserving. Furthermore, the basis of the logical space has been chosen such that CSP physical Pauli operators always implement logical Pauli gates, so each P^i is a dressed CSP operator supported on R^c implementing $VP_iV^\dagger = P_L^i \otimes P_J^i$. Since A implements $P_L \otimes \mathbb{1}$, we have $VAJ^\dagger = \sum_i \lambda_i P_L^i \otimes P_J^i = P_L \otimes \mathbb{1}_J$. Since the Pauli gates form a basis of logical operators, at least one P^i is a bare CSP operator implementing $P_L \otimes \mathbb{1}_J$, completing the proof.

($\impliedby$) Any logical operator A_L can be decomposed into Pauli operators as $\sum_i \lambda_i P_L^i$. Then each of these Pauli operators can be implemented as a bare CSP operator P^i on R^c . Then $\sum_i \lambda_i P^i$ is a bare CSP operator supported on R^c implementing A_L , concluding the proof. ■

Similar proofs can be constructed for the bare-trivial and dressed-trivial properties; however, the reasoning breaks down for the dressed-cleanable property. In the converse direction of the above proof, it was crucial that the linear combination of bare CSP operators $\sum_i \lambda_i P^i$ was also a bare CSP operator. However, this does not apply to dressed CSP operators, which are not closed under linear combination. Thus, we need to define two separate notions of dressed cleanability for stabilizer subsystem codes—one based on general operators and one based on Pauli operators. The latter of these is the one that was used in the main text and referred to simply as “dressed-cleanable”—because it is the one that is useful for the Pastawski-Yoshida lemma, Lemma II.18.

Definition A.3. A region R of a subsystem stabilizer code is defined to be

- (a) *general dressed-cleanable* if and only if for any logical operator A_L there exists a dressed CSP operator A supported on R^c that implements $A_L \otimes A_J$;
- (b) *Pauli dressed-cleanable* if and only if for any logical Pauli operator P_L there exists a dressed CSP operator P supported on R^c that implements $P_L \otimes P_J$.

It can still be shown analogously to the first part of the proof of Lemma A.2 that general dressed-cleanable regions are also Pauli dressed-cleanable regions.

The one other difference for stabilizer subsystem codes is that the converse directions for Lemmas II.3 and II.14 are also true (see Ref. [54]), as mentioned in the main text. These are shown specifically for Pauli dressed cleanability and not general dressed cleanability. The web of relationships for stabilizer subsystem codes is shown in Fig. 10.

To the best of our knowledge, this is the extent of the known relationships between different correctability properties in both subsystem and stabilizer subsystem codes. This leaves many open questions: Are dressed-trivial and bare-cleanable equivalent for any subsystem code? Is the general dressed-cleanable property of stabilizer subsystem codes equivalent to Pauli dressed cleanability after all, is it equivalent to bare cleanability, or is it distinct from both? In HaPPY codes, such as in Fig. 7, regions are found that are Pauli dressed-cleanable and not bare-cleanable, which already shows that some of the above-mentioned properties must be distinct. Perhaps similar counterexamples can be found demonstrating that the relationships shown here are all that can be generally expressed without additional information. We leave this to future work.

APPENDIX B: HOMOGENEOUS PLANAR HYPERBOLIC HaPPY CODES

In Ref. [74], a range of possible HaPPY codes were considered by systematically exploring possible homogeneous hyperbolic tilings characterized by Schläfli symbols. Schläfli symbols are a characterization of a regular polygonal tiling via a pair of integers n, k , such that k n -sided polygons (n -gons) meet at each vertex. Such a tiling is hyperbolic if for each polygon, the sum of the internal angles exceeds $(n - 2)\pi$, which is equivalent to

$$\frac{1}{n} + \frac{1}{k} \leq \frac{1}{2}. \quad (\text{B1})$$

In Ref. [74], the perfect tensors of HaPPY are associated with polygons of this tiling; however, for our purposes we consider the dual identification of these tensors with the *vertices* of the graph instead. Thus, for a network with

FIG. 10. Relationships between properties of physical regions in subsystem codes. The blue relationship holds for regions with complementary recovery, such as those in holographic codes.

a rank- $2N$ perfect tensor, we should have $k = 2N - 1$ as the number of in-plane legs. These two conventions are related by Poincaré duality, where faces and vertices are interchanged, which corresponds to switching $n \leftrightarrow k$ in the Schläfli symbol. The symmetry of Eq. (B1) shows that a tiling is hyperbolic if and only if its dual is hyperbolic, as expected.

Of course, the HaPPY tensor network is more than just a hyperbolic tiling filled with a network of perfect tensors; it must also be an isometry from the bulk to the boundary. The former does not imply the latter, as one can see from the case where $k = 3$. In this case, a local bulk operator cannot be pushed out to the boundary; whenever an operator is pushed to a new vertex, there are two inputs (one from the leg “pushed” from the previous vertex, and one from the bulk input), so the result of the pushing is that some operator is pushed to both remaining legs of the tensor. In this way, there is no control over the pushing process to guide the output toward the boundary; it inevitably loops back on itself as soon as given the chance, because it follows all available paths. Thus, an isometry does not occur for this case.

Whenever $k > 3$, it appears that the tensor network does result in an isometry. In Ref. [74], these cases were studied separately according to $n = 3$ and $n > 3$. A central perfect tensor is placed to begin with, and then consecutive layers of polygons are added. This is done such that each vertex is on the boundary of the tiling when it is in the most recently added layer, and it is in the interior afterwards.

Then the vertices of the network can be characterized according to how many other vertices they are connected to from the same layer or previous layers. Each such vertex is connected to exactly two from the same layer as itself, as it is on the boundary of the tiling. It can connect to no vertices from previous layers (only two of those from the same layer as itself), to one vertex from a previous layer (and thus a total of three, including those of the same layer), or to two vertices from previous layers (for a total of four vertices). We label these types of vertices as a , b , and c , respectively, and note that the third type, c , appears only when $n = 3$. The initial vertex effectively counts as a distinct category, i , which connects to no vertices from the same layer or previous layers.

The construction of each new layer can be surmised using replacement rules of these vertex types as follows [74,75]:

$$n = 3 : \begin{cases} a \rightarrow cb^{k-4}, \\ b \rightarrow cb^{k-5}, \\ c \rightarrow cb^{k-6}, \end{cases} \quad (\text{B2})$$

$$n > 3 : \begin{cases} a \rightarrow a^{n-4}b(a^{n-3}b)^{k-3}, \\ b \rightarrow a^{n-4}b(a^{n-3}b)^{k-4}. \end{cases} \quad (\text{B3})$$

The resultants of these replacement rules are written in a specific order to denote a particular ordering of each layer (e.g., one can use a clockwise convention such that the new vertices are in the appropriate order from left to right). We also require the rules for i -type vertices, which were not present in Ref. [74] due to the different convention, which are as follows for all n :

$$i \rightarrow (ba^{n-3})^k. \quad (\text{B4})$$

Our proof strategy is as follows: We construct the network such that the initial site is the one onto which a local bulk operator is applied. This operator is then pushed through to three different choices of region S_1 , S_2 , and S_3 . Each of these regions can fully reconstruct an arbitrary bulk logical operator on the initial site, so their complements are correctable, and if the union of these complements covers the full boundary space, we can apply Lemma II.18 to obtain the desired result. Thus, we must show that S_1 , S_2 , and S_3 can be constructed such that no site of the boundary lies in the intersection of all three regions. We proceed by partitioning the boundary into $k = 2N - 1$ sets, $\{M_j\}_{j=1}^{j=2N-1}$, each of which corresponds to some b_j , one of the b -type vertices that directly connects to the initial vertex. Each of these vertices comes with $n - 3$ a -type vertices from the i replacement rule, Eq. (B4). Together, these vertices, along with all of the vertices obtained from them via repeated application of the replacement rules in Eq. (B3), form the set M_j .

One can push out the bulk operator acting on the initial vertex to any choice of N of the b_j vertices. Consider what would happen if one could construct a pushing protocol such that an operator acting on each b_j can always be pushed out into just the boundary region M_j . Then the initial bulk operator can be pushed to any N of the M_j , and we could use these sets to form S_1 , S_2 , and S_3 . For example, let $S_1 = \bigcup_{j=1}^{j=N} M_j$, $S_2 = \bigcup_{j=N+1}^{j=2N-1} M_j$, and $S_3 = \bigcup_{j=N+1}^{j=2N-1} M_j$. Each S_i consists of the union of N of the M_j regions, and no M_j belongs to all three, so these choices of S_i are appropriate choices that allow application of the main lemma, Lemma II.18, and thus no non-Clifford gates can be applied transversally.

Definition B.1 (replacement pushing assumption). *An operator acting on a vertex b_j connected to the initial vertex can be pushed out to M_j , which is defined as the image of b_j under repeated applications of the replacement rules Eq. (B3).*

In other words, the replacement pushing assumption means that we can partition the boundary into $2N - 1$ regions $\{R_j\}$ such that each $b_j \in \mathcal{E}[R_j]$, with the entanglement wedge $\mathcal{E}$ as defined in Definition II.4.

When does the replacement pushing assumption apply? The replacement rules are constructed such that whenever

they are applied to an existing vertex and a new b or c vertex is obtained, it must connect to that existing vertex. Thus, if at least N vertices of b or c type are obtained from any replacement rule, then the replacement pushing assumption applies straightforwardly; each vertex can be pushed into its replacement, which is repeated until the boundary is reached.

First, let us consider the $n > 3$ case. The rules for a and b result in $k - 2$ and $k - 3$ new b -type vertices, respectively. Thus, so long as $k - 3 \geq N$ (i.e., $k \geq 7$), replacement pushing applies, and the result follows.

For $n = 3$, because none of the replacement rules in Eqs. (B3) and (B4) produce any a vertices, we need to consider only b and c (again, this is different to Ref. [74], where the conventions necessitate starting with some a vertices). Thus, if $N \geq k - 5$ (i.e., $k \geq 11$), replacement pushing applies, and the result follows.

The cases of $\{n, k\}$ that produce valid HaPPY networks (i.e., have odd $k = 2N - 1$, satisfy Eq. (B1), and form actual isometries, i.e., $k \neq 3$) but for which replacement pushing does not straightforwardly apply according to the above arguments, are (1) $\{n > 3, 5\}$, (2) $\{3, 9\}$, and (3) $\{3, 7\}$. In each of these cases, replacement pushing does still apply, but with a little more effort required to complete the argument.

In case 1, replacement pushing does straightforwardly apply for a -type vertices, as their replacement rule gives three b -type vertices, and $N = (k + 1/2) = 3$. However, b -type vertices are potential causes for concern, as operators acting here must be pushed to $N = 3$ connecting vertices, but there are only two available connected b -type vertices in the succeeding layer. Thus, one leg must be pushed into *along* the boundary (i.e., there must be some pushing to the same layer). However, this is fine; we just push in the direction of the nearest a -type vertex that belongs to the same vertex set M_j —noting that because $n > 3$, there always exists at least one such a -type vertex. Even if the b -type vertex in question does not connect to an a -type vertex directly, it can push to a neighboring b -type vertex, which will eventually connect to an appropriate a -type vertex from the same set M_j .

In case 2, each b -type vertex is mapped under replacement to a $cbbb$ string of vertices, so these straightforwardly obey the replacement pushing assumption. Each c -type vertex is mapped under replacement with $cbbb$, but we must push to $N = 5$ vertices. Because of the replacement rules, no two c -type vertices are ever adjacent, so there must be a b -type vertex to the right of each c -type vertex. Furthermore, this b -type vertex must belong to the same set M_j , again because of the nature of the replacement rules. Thus, we simply push from each c -type vertex onto its four descendant vertices, and for the remaining leg, push into the b -type vertex to its right—a protocol that satisfies the replacement pushing assumption.

Case 3 is slightly more complicated as we must use different pushing rules for b -type vertices depending on their position. Since the $n = 3$ replacement rules always generate a c for each vertex, and c -type vertices by definition connect with *two* vertices from the previous layer, the c in a given vertex's replacement is also connected to that of the vertex to its left. Thus, a vertex has three options for pushing: it can push "forward" to the vertices in its replacement, it can push "sideways" to adjacent vertices of the same layer, or it can push "diagonally" to the c -type vertex in the replacement of its neighbor on the right. Each b -type vertex is mapped under replacement to a cb string of vertices, but must be pushed along $N = 4$ legs. For any b -type vertex that is not on the very far right of the corresponding set M_j , it can be pushed into the three vertices in its replacement, and also to the c -type vertex in the replacement of its neighbor on the right. The rightmost vertex, which is always a b -type vertex, can be pushed to its left neighbor (which is always a b -type vertex), as well as to the three in its replacement. The rightmost vertex must always have a b -type neighbor because of the replacement rules; although some b -type vertices have c -type vertices on either side, these are never on the boundary of the M_j set.

Finally, each c -type vertex is mapped under replacement to cb , but must also be pushed along four legs. Again, we can push into the two in its replacement, as well as the c -type vertex in the replacement of its rightward neighbor. For the final leg, we can push sideways to an adjacent b -type vertex from the same M_j .

None of these pushing rules are in conflict with one another, so we have shown a consistent pushing protocol that keeps each operator in the same M_j (i.e., the replacement pushing assumption is satisfied).

All of these ideas still apply in the case of perfect planar and block perfect tensors as well, as we never need to push to nonadjacent legs. Similar arguments can be generalized to nonhomogeneous codes (such as in Fig. 6), as well as higher-dimensional HaPPY codes.

---

[1] J. M. Maldacena, The large N limit of superconformal field theories and supergravity, *Int. J. Theor. Phys.* **38**, 1113 (1997).
 [2] E. Witten, Anti de sitter space and holography, *Adv. Theor. Math. Phys.* **2**, 253 (1998).
 [3] G. t. Hooft, in *Conference on Highlights of Particle and Condensed Matter Physics* (1993), Vol. 930308, p. 284.
 [4] L. Susskind, The world as a hologram, *J. Math. Phys.* **36**, 6377 (1995).
 [5] A. Almheiri, X. Dong, and D. Harlow, Bulk locality and quantum error correction in AdS/CFT, *J. High Energ. Phys.* **2015**, 163 (2015).
 [6] F. Pastawski, B. Yoshida, D. Harlow, and J. Preskill, Holographic quantum error-correcting codes: Toy models for the

bulk/boundary correspondence, *J. High Energ. Phys.* **2015**, 149 (2015).
 [7] P. Hayden, S. Nezami, X.-L. Qi, N. Thomas, M. Walter, and Z. Yang, Holographic duality from random tensor networks, *J. High Energ. Phys.* **2016**, 9 (2016).
 [8] W. Donnelly, B. Michel, D. Marolf, and J. Wien, Living on the edge: A toy model for holographic reconstruction of algebras with centers, *J. High Energ. Phys.* **2017**, 93 (2017).
 [9] C. Cao and B. Lackey, [arXiv:2010.05960 \[hep-th, physics:quant-ph\]](#) (2020).
 [10] T. Farrelly, R. J. Harris, N. A. McMahon, and T. M. Stace, [arXiv:2012.07317 \[quant-ph\]](#) (2020).
 [11] R. J. Harris, N. A. McMahon, G. K. Brennen, and T. M. Stace, Calderbank-shor-steane holographic quantum error-correcting codes, *Phys. Rev. A* **98**, 052301 (2018).
 [12] R. J. Harris, E. Coupe, N. A. McMahon, G. K. Brennen, and T. M. Stace, [arXiv:2008.10206 \[quant-ph\]](#) (2020).
 [13] C. Cao, J. Pollack, and Y. Wang, [arXiv:2103.08631](#) (2021).
 [14] A. Jahn and J. Eisert, Quantum Science and Technology (2021).
 [15] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, *Phys. Rev. A* **52**, R2493 (1995).
 [16] D. Aharonov and M. Ben-Or, Fault-tolerant quantum computation with constant error Rate, *SIAM J. Comput.* **38**, 1207 (2008).
 [17] J. Preskill, in *Proceedings of 37th Conference on Foundations of Computer Science* (WORLD SCIENTIFIC, 1997), p. 56, [arXiv:9712048 \[quant-ph\]](#).
 [18] F. J. Wegner, Duality in generalized ising models and phase transitions without local order Parameters, *J. Math. Phys.* **12**, 2259 (1971).
 [19] R. B. Laughlin, Anomalous Quantum Hall Effect: An Incompressible Quantum Fluid with Fractionally Charged Excitations, *Phys. Rev. Lett.* **50**, 1395 (1983).
 [20] X. G. WEN, Topological orders in rigid states, *Int. J. Mod. Phys. B* **04**, 239 (1990).
 [21] A. Y. Kitaev, Fault-tolerant quantum computation by anyons, *Ann. Phys. (N. Y.)* **303**, 2 (2003).
 [22] S. B. Bravyi and A. Y. Kitaev, [arXiv:9811052 \[quant-ph\]](#) (1998).
 [23] M. H. Freedman and D. A. Meyer, Projective plane and planar quantum Codes, *Found. Comput. Math.* **1**, 325 (2001).
 [24] E. Knill, [arXiv:0404104 \[quant-ph\]](#) (2004).
 [25] S. Bravyi and A. Kitaev, Universal quantum computation with ideal clifford gates and noisy ancillas, *Phys. Rev. A - At. Mol. Opt. Phys.* **71**, 022316 (2005).
 [26] E. Knill and R. Laflamme, [arXiv:9608012 \[quant-ph\]](#) (1996).
 [27] B. Eastin and E. Knill, Restrictions on Transversal Encoded Quantum Gate Sets, *Phys. Rev. Lett.* **102**, 110502 (2009).
 [28] S. Bravyi and R. Konig, Classification of Topologically Protected Gates for Local Stabilizer Codes, *Phys. Rev. Lett.* **110**, 170503 (2013).
 [29] M. E. Beverland, O. Buerschaper, R. Koenig, F. Pastawski, J. Preskill, and S. Sijher, Protected gates for topological quantum field theories, *J. Math. Phys.* **57**, 022201 (2016).
 [30] This does not include the more general class of gates and codes with boundaries considered in Refs. [76,77].

- [31] F. Pastawski and B. Yoshida, Fault-tolerant logical gates in quantum error-correcting codes, *Phys. Rev. A* **91**, 012305 (2015).
- [32] P. Webster and S. D. Bartlett, Locality-preserving logical operators in topological stabilizer codes, *Phys. Rev. A* **97**, 012330 (2018).
- [33] H. Bombín, Gauge color codes: Optimal transversal gates and gauge fixing in topological stabilizer codes, *New J. Phys.* **17**, 083002 (2015).
- [34] H. Bombin and M. A. Martin-Delgado, Topological Computation Without Braiding, *Phys. Rev. Lett.* **98**, 160502 (2007).
- [35] D. Harlow and H. Ooguri, [arXiv:1810.05338 \[gr-qc, physics:hep-lat, physics:hep-ph, physics:hep-th\]](#) (2019).
- [36] D. Harlow and H. Ooguri, Constraints on Symmetries from Holography, *Phys. Rev. Lett.* **122**, 191601 (2019).
- [37] Despite the no-go result in Ref. [35], such global symmetries are possible when domain walls of the boundary symmetry cannot be locally mapped back into the code subspace. This does not contradict Ref. [35] as it violates the assumptions made there, and we plan to discuss this further in forthcoming work [78].
- [38] D. Harlow, The Ryu-Takayanagi formula from quantum error correction, *Commun. Math. Phys.* **354**, 865 (2017), [arXiv:1607.03901](#).
- [39] F. Pastawski and J. Preskill, Code Properties from Holographic Geometries, *Phys. Rev. X* **7**, 021022 (2017).
- [40] H. Kamal and G. Penington, [arXiv:1912.02240](#) (2019).
- [41] Operations acting on the gauge subsystem that act trivially on the logical subsystem are unimportant, as we do not care about the state of the gauge system. In this way, they are similar to gauge transformations in gauge theory, hence the name.
- [42] One can generalize this definition by considering correctability with respect to any particular noise channel, rather than just erasure channels as we do here. However, we are primarily interested in the implications of the geometric local structure of holographic codes, which is most readily apparent when studying correctability with respect to erasure of particular regions.
- [43] S. Bravyi and B. Terhal, A no-go theorem for a two-dimensional self-correcting quantum memory based on stabilizer codes, *New J. Phys.* **11**, 043029 (2009).
- [44] This follows from the fact that any operator A supported on the Hilbert space associated with $S_1 \cup S_2$ can be written as a linear combination of product operators on the spaces associated with S_1 and $S_2 \setminus S_1$, $A = \sum_i A_1^i \otimes A_{2 \setminus 1}^i$, with an implicit identity acting on systems not labeled. By correctability of R with respect to $S_{1,2}$, each operator A_1^i and $A_{2 \setminus 1}^i$ can be implemented by a CSP operator supported on R^c . These operators can be multiplied and linearly combined appropriately to implement A with a CSP operator on R^c , and thus R is correctable with respect to $S_1 \cup S_2$.
- [45] We conjecture that our definition coincides with that of the greedy entanglement wedge from HaPPY—essentially, perfect tensors are already so highly constrained that the entire algebra of operators on a subsystem can be correctable only if this is already guaranteed by the properties of perfect tensors.
- [46] More precisely we mean that for a local encoded subsystem with a fixed price, a code that satisfies complementary recovery leads to the maximum possible distance for that subsystem.
- [47] The size of this residual region for a connected boundary component was argued to be constant in Ref. [6].
- [48] There can be discontinuities in the entanglement wedge map, such as when a phase transition occurs between connected and disconnected entanglement wedges. These do not contradict our definition of smoothness.
- [49] Recall that the center $\mathcal{Z}[A]$ of an algebra A is the set of elements of A that commute with all of A .
- [50] In other words, for $j = 1, \dots, d-1$, $\omega_d^j \notin \mathcal{S}$. In the qubit case, this condition coincides with the familiar $-1 \notin \mathcal{S}$.
- [51] The tensor product structure and local basis for the encoded space $\mathcal{H} = \mathcal{H}_L \otimes \mathcal{H}_J$ are unique only up to transformation by a Clifford operation, which preserves the encoded Pauli group.
- [52] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, 2010).
- [53] One can equivalently define a stabilizer subsystem code by choosing some generally non-Abelian subgroup $\mathcal{G} \subset \mathcal{P}$, defining the stabilizer group $\mathcal{S}$ as the center of $\mathcal{G}$, and then defining the division into logical and junk subsystems such that elements of the quotient group $\mathcal{G}/\mathcal{S}$ are identified with dressed CSP operators that act only on the junk subsystem [79,80].
- [54] S. Bravyi, Subsystem codes with spatially local generators, *Phys. Rev. A* **83**, 012320 (2011).
- [55] T. Jochym-O'Connor, A. Kubica, and T. J. Yoder, The Disjointness of Stabilizer Codes and Limitations on Fault-Tolerant Logical Gates, *Phys. Rev. X* **8**, 021047 (2018).
- [56] P. Faist, S. Nezami, V. V. Albert, G. Salton, F. Pastawski, P. Hayden, and J. Preskill, Continuous Symmetries and Approximate Quantum Error Correction, *Phys. Rev. X* **10**, 041018 (2020).
- [57] M. P. Woods and Á. M. Alhambra, Continuous groups of transversal gates for quantum error correcting codes from finite clock reference frames, *Quantum* **4**, 245 (2020).
- [58] S. Burton and D. Browne, [arXiv:2012.05842](#) (2020).
- [59] This differs slightly from the standard definition of the Clifford hierarchy [28], which consists of unitaries such that $UPU^\dagger \in \mathcal{C}_{i-1}$. It was shown in Ref. [31] that the definitions coincide for $i > 1$. For $i = 1$, this definition gives $\mathcal{C}_1 = \mathcal{P} \times \mathbb{C}$, whereas the standard definition gives $\mathcal{C}_1 = \mathcal{P}$. This definition also allows $\mathcal{C}_0$ to be nonempty, unlike the standard definition.
- [60] In particular, a logarithmic distance and a nonzero error threshold.
- [61] This is a specific instance of global symmetry inherited from local symmetry, which is common in the tensor network literature [81–86]. An interesting open question is whether a converse statement of global implying local tensor symmetry can be proven for holographic tensor networks.
- [62] T. J. Yoder, R. Takagi, and I. L. Chuang, Universal Fault-Tolerant Gates on Concatenated Stabilizer Codes, *Phys. Rev. X* **6**, 031039 (2016).

- [63] J. Berger and T. J. Osborne, Perfect tangles, [arXiv:1804.03199 \[quant-ph\]](#) (2018).
- [64] The random tensor networks in Ref. [7] provide nonstabilizer examples.
- [65] We use $d_{(c)}$ to denote d if R is not connected, and d_c if it is.
- [66] Unlike in Sec. IV, the d_c variant is not strictly stronger than the d variant. For example, a code may exist where only disconnected regions satisfy complementary recovery.
- [67] One can also express this in terms of the bulk rather than the boundary, i.e., by requiring $\mathcal{E}[R^c]^c \subset B(\mathcal{E}[R], \delta)$, and then enforcing $\kappa_R(\delta) \leq \varepsilon$. This has the nice physical interpretation that δ bounds the thickness of the residual region contained in neither $\mathcal{E}[R]$ nor $\mathcal{E}[R^c]$. For example, connected regions of HaPPY codes satisfy it with $\delta = 1$; however, formulating it in terms of the boundary makes Theorem IV.5 slightly stronger.
- [68] We expect $B(R_i, s)$ to scale as $|\partial R_i|s$ for small s , so a region with small boundary-to-bulk ratio, such as a ball in flat space, suffices.
- [69] V. E. Hubeny, Extremal surfaces as bulk probes in AdS/CFT, [J. High Energ. Phys. **2012**, 93 \(2012\)](#).
- [70] C. Bény, A. Kempf, and D. W. Kribs, Generalization of Quantum Error Correction via the Heisenberg Picture, [Phys. Rev. Lett. **98**, 100502 \(2007\)](#).
- [71] C. Bény, A. Kempf, and D. W. Kribs, Quantum error correction of observables, [Phys. Rev. A **76**, 042303 \(2007\)](#).
- [72] A. M. Turing, Finite approximations to lie Groups, [Ann. Math. **39**, 105 \(1938\)](#).
- [73] P. Faist, S. Nezami, V. V. Albert, G. Salton, F. Pastawski, P. Hayden, and J. Preskill, Perfect Tangles, [Phys. Rev. X **10**, 041018 \(2020\)](#).
- [74] A. Jahn, Z. Zimborás, and J. Eisert, Central charges of aperiodic holographic tensor-network models, [Phys. Rev. A **102**, 042407 \(2020\)](#).
- [75] L. Boyle, M. Dickens, and F. Flicker, Conformal Quasicrystals and Holography, [Phys. Rev. X **10**, 011009 \(2020\)](#).
- [76] G. Zhu, M. Hafezi, and M. Barkeshli, Quantum origami: Transversal gates for quantum computation and measurement of topological order, [arXiv:1711.05752](#) (2017).
- [77] A. Lavasani, G. Zhu, and M. Barkeshli, Universal logical gates with constant overhead: instantaneous Dehn twists for hyperbolic quantum codes, [Quantum **3**, 180 \(2019\)](#).
- [78] V. Calvera, S. Cree, K. Dolev, and D. J. Williamson, Gauging the bulk: generalized gauging maps and holographic codes, [arXiv:2108.11402](#) (2021).
- [79] D. Poulin, Stabilizer Formalism for Operator Quantum Error Correction, [Phys. Rev. Lett. **95**, 230504 \(2005\)](#).
- [80] D. Bacon, Operator quantum error-correcting subsystems for self-correcting quantum memories, [Phys. Rev. A - At. Mol. Opt. Phys. **73**, 012340 \(2006\)](#).
- [81] D. Pérez-García, M. Sanz, C. E. González-Guillén, M. M. Wolf, and J. I. Cirac, Characterizing symmetries in a projected entangled pair state, [New J. Phys. **12**, 25010 \(2010\)](#).
- [82] S. Singh, R. N. C. Pfeifer, and G. Vidal, Tensor network decompositions in the presence of a global symmetry, [Phys. Rev. A - At. Mol. Opt. Phys. **82**, 050301\(R\) \(2010\)](#).
- [83] S. Singh and G. Vidal, Symmetry-protected entanglement renormalization, [Phys. Rev. B - Condens. Matter and Mater. Phys. **88**, 121108\(R\) \(2013\)](#).
- [84] D. J. Williamson, N. Bultinck, M. Mariën, M. B. Sahinoglu, J. Haegeman, and F. Verstraete, Matrix product operators for symmetry-protected topological phases: Gauging and edge theories, [Phys. Rev. B **94**, 205150 \(2016\)](#).
- [85] J. C. Bridgeman and D. J. Williamson, Anomalies and entanglement renormalization, [Phys. Rev. B **96**, 125104 \(2017\)](#).
- [86] D. J. Williamson, N. Bultinck, and F. Verstraete, [arXiv:1711.07982](#) (2017).