Approximate *t*-Designs in Generic Circuit Architectures

Daniel Belkin,[1,†] James Allen,[1,†] Soumik Ghosh,[2] Christopher Kang,[2] Sophia Lin,[2] James Sud,[2] Frederic T. Chong,[2] Bill Fefferman,[2] and Bryan K. Clark[1,*]

[1] *Institute for Condensed Matter Theory and IQUIST and NCSA Center for Artificial Intelligence Innovation and Department of Physics, University of Illinois at Urbana-Champaign, Illinois 61801, USA*

[2] *Department of Computer Science, The University of Chicago*

(Received 5 April 2024; accepted 25 October 2024; published 17 December 2024)

Unitary *t*-designs are distributions on the unitary group whose first *t* moments appear maximally random. Previous work has established several upper bounds on the depths at which certain specific random quantum circuit ensembles approximate *t*-designs. Here we show that these bounds can be extended to any fixed architecture of Haar-random two-site gates. This is accomplished by relating the spectral gaps of such architectures to those of one-dimensional brickwork architectures. Our bound depends on the details of the architecture only via the typical number of layers needed for a block of the circuit to form a connected graph over the sites. When this quantity is bounded, the circuit forms an approximate *t*-design in at most linear depth. We give numerical evidence for a stronger bound that depends only on the number of connected blocks into which the architecture can be divided. We also give an implicit bound for nondeterministic architectures in terms of properties of the corresponding distribution over fixed architectures.

DOI: 10.1103/PRXQuantum.5.040344

I. INTRODUCTION

Random quantum circuits are an important tool in the study of natural and engineered quantum systems. In quantum computing, random circuits have been suggested for randomized benchmarking [1,2], security, and state preparation [3]. Recent claims of quantum supremacy have hinged on the hardness of classical simulation of random circuits [4–9]. Random circuits have also been proposed as models for information scrambling in black holes [10,11], and more general random tensor networks have been used as an explicit construction of the holographic duality in AdS/CFT [12]. In quantum information theory, random circuits are the standard setting in which to study measurement-induced phase transitions [13], and they are used as an analytically tractable model of quantum ergodicity and chaos [14]. Random circuit models also serve as an interesting theoretically tractable model for more complicated realistic physical systems. Their maximally generic dynamics are often a valuable source of intuition as to what one should expect under typical time evolution.

The geometric structure of a quantum circuit plays a critical role in the flow of information. While initial work on quantum computing focused on one-dimensional (1D) architectures, 2D layouts, such as the Sycamore processor used in the quantum advantage experiment of Ref. [15], have become increasingly popular. Recent work has explored modular architectures, in which fully connected nodes are sparsely connected with each other [16]. Circuit models for physical systems often require geometric locality on some two- or three-dimensional lattice. In certain cases, such as the Sachdev–Ye–Kitaev model, the interactions are instead all to all. Meanwhile, circuit architectures with the connectivity pattern of a tree or the multiscale entanglement renormalization ansatz [17] are a natural setting in which to study holography. Such circuits may also be useful for robust quantum simulations of many-body systems [18].

In the limit of large depth, a sufficiently well-connected random quantum circuit will eventually scramble quantum information [19]. Perhaps the most basic question about random quantum circuits is the rate of this scrambling. This is quantified by the *approximate t-design depth* [20], which captures the depth at which an observer with access to at most *t* measurements can no longer reliably distinguish the circuit from a random global unitary.

*Contact author: bkclark@illinois.edu

†Co-first authors.

Prior work has given bounds on the t-design depths for a few special classes of architectures. However, it was not previously clear how the rate of information scrambling depended on the spatial structure of the circuit. In particular, one might have expected irregular or modular architectures to give qualitatively different behavior, e.g., exponentially slow convergence. The main goal of this work is to show that any reasonable architecture forms an approximate t-design in linear depth. We bound the rate of convergence in terms of the connectedness of the architecture.

A. Prior work

The Haar measure on the unitary group is clearly a fixed point of any quantum circuit distribution, since it is invariant under any unitary gate. Emerson *et al.* [19] showed that random circuits satisfying a universality condition converge to this fixed point in the limit of large depth. But this uniform convergence is very slow, requiring a circuit depth that scales exponentially with system size N. On the other hand, the expected values of specific observables sometimes approach their Haar values much faster (e.g., in depth $\log N$) [21]. But this fast convergence depends on specific details of the observables considered and is not necessarily universal for other quantities of interest. The approximate t-design depth was first introduced in Ref. [20] as an intermediate measure of convergence. It is strong enough to guarantee convergence of any experimentally observable property, but occurs much faster than the uniform convergence of measure.

Much of the prior work on approximate t-design depths has focused on the 1D brickwork architecture. Brandão *et al.* [22] showed that the approximate t-design depth in this case was at most $O(t^{9.5+o(1)}N)$. For local Hilbert space dimension $q = 2$, Haferkamp [23] tightened this bound to $O(t^{5+o(1)}N)$. Hunter-Jones [24] gave a mapping to a statistical-mechanical model of interacting domain walls and used it to establish tighter bounds when either $t = 2$ or $q \to \infty$. Harrow and Mehraban [25] extended this work to a particular family of D-dimensional brickwork architectures. In the limit of small ϵ and large N, they established that the approximate t-design depth scales as at most $O(N^{1/D})$ (although the dependence on t remains an open question). Schuster *et al.* [26] and LaRacuente and Leditzky [27], meanwhile, constructed families of architectures for which the approximate t-design depth is $O[\mathrm{poly}(t)\log N]$.

The other class of prior work focuses on what we term *nondeterministic architectures*, in which the spatial structure of the architecture is also random. Typically, gate locations are assumed to be drawn independently and identically from the uniform distribution over the edges of some graph over the sites. In this context, Ambainis and Emerson [20] established an approximate 2-design size of at most $O(N^2)$ gates for the all-to-all graph. Brandão *et al.* [22] found $O(t^{9.5}N^2)$ for the linear graph, which Oszmaniec *et al.* [28] extended to $O[t^{9.5}\log^4(t)N^3]$ for any graph that admits a Hamiltonian path. Mittal and Hunter-Jones [29] developed an alternative strategy that yields a bound of the form $N^{O(\log N)}\mathrm{poly}(t)$ for arbitrary graphs. In addition, they gave a bound of the form $O[|E|N\mathrm{poly}(t)]$ for graphs with $|E|$ edges, bounded degree, and bounded effective spanning-tree height. For certain bounded t and degree, the requirement of bounded spanning-tree height can be relaxed. The results of Ref. [30], meanwhile, imply a bound of

$$O[N^9 t(\log t)^3(\log N + \log\log t)]$$

gates for arbitrary graphs.

B. Summary of results

We obtain bounds on the t-design depth for all architectures, with stronger bounds if the architecture satisfies certain properties. First, consider a periodic architecture composed of ℓ complete layers of two-site gates on N qubits. For large t, our bound for the ϵ-approximate t-design depth becomes

$$d_* = t^{[5+o(1)](\ell-1)}\left(2Nt\log 2 + \log\frac{1}{\epsilon}\right). \tag{1}$$

We can relax each of these assumptions to obtain looser bounds for larger classes of circuits. The most general result covers an architecture with local Hilbert space dimension q that may not be periodic or consist of complete layers [31]. We partition such a circuit into blocks of layers such that the gates in each block form a connected graph over all the sites. Let $\bar{\ell}$ be the average number of layers per block [32]. For these architectures, our bound is of the form

$$d \geq t^{[15.2+o(1)](\bar{\ell}-1)}\left(2Nt\log q + \log\frac{1}{\epsilon}\right). \tag{2}$$

We also give numerical evidence for two conjectures under which Eq. (2) can be strengthened to

$$d_* = \frac{[2Nt\log q + \log(1/\epsilon)]\bar{\ell}}{2\log[(q^2+1)/2q]}. \tag{3}$$

For nondeterministic architectures, we obtain an implicit bound in terms of the joint distribution of the effective number of layers in the circuit and the number of connected blocks.

C. Structure of the proof

We wish to prove that the distribution induced by a random circuit architecture approaches the Haar measure.

Following previous work on approximate t-designs [33], we begin by expressing the frame potential as a tensor network of single-gate moment operators (Sec. II). In Sec. III, we focus on the case of periodic complete architectures and show that the t-design depth is determined by the spectral gap of the transfer matrix T. (The assumptions of completeness and periodicity will be relaxed in Secs. VII A and VII B, respectively.)

We then decompose the transfer matrix into a product of ℓ layers of gates. Each layer is an orthogonal projection operator. We wish to bound the spectral gap of the transfer matrix in terms of the geometry of the subspaces to which the layers project. The key insight at this stage is to consider the way adding a new layer shrinks the unit eigenspace of the product, reducing the norm of the excluded vectors. In Sec. IV, we bound the spectral gap in terms of these norm reductions. The impact of a new layer on the unit eigenspace can be represented by a graph of nodes and edges, which we term the *cluster-merging picture*.

The next step is to simplify the cluster-merging picture at each layer. In Sec. V, we show that you can unravel the cluster-merging graph of each layer into a collection of loops without increasing the spectral gap. This is useful because each loop is the transfer matrix of a periodic 1D brickwork architecture. At this stage we have obtained a lower bound on our transfer matrix spectral gap in terms of the spectral gaps of 1D brickwork architectures.

In Sec. VI, we proceed to show that the 1D brickwork spectral gap itself may be bounded by previous results on 1D approximate t-design depths. Many of these results were actually originally proven in terms of the spectral gap, but our argument applies even to bounds obtained by other methods. Together these steps allow us to turn a bound on the 1D t-design depth into a bound for generic architectures.

We also discuss extensions of our techniques to other architectures. We show how our bounds can be applied in expectation to architectures in which gate locations are drawn randomly, giving an implicit bound in terms of properties of the induced distribution over fixed architectures. We observe that our techniques may be adapted to give tighter bounds for architectures with special structure and give an explicit example for the case of higher-dimensional brickwork architectures.

D. Definitions

A quantum circuit on N sites of local dimension q corresponds to a unitary $U_c \in U(q^N)$. A random quantum circuit architecture then induces a measure ε_C on $U(q^N)$. Define the associated t-fold channel

$$\Phi_{\mathrm{RQC}}(\rho) = \int_{\varepsilon_C} U_C^{\otimes t} \rho (U_C^\dagger)^{\otimes t} \mathrm{d}U_C. \tag{4}$$

Definition 1. An ϵ-approximate unitary t-design [22] is a measure ε_C on $U(q^N)$ such that the diamond norm distance between the corresponding t-fold channel Φ_{RQC} and that of the Haar measure is at most ϵ:

$$\|\Phi_{\mathrm{RQC}} - \Phi_{\mathrm{Haar}}\|_\diamond \leq \epsilon. \tag{5}$$

We often shorten "ϵ-approximate unitary t-design" to "(ϵ, t)-design."

Definition 2. We call an L-layer random circuit architecture on N sites *complete* if each of the L layers consists of $N/2$ Haar-random two-site unitary gates. In other words, exactly one gate acts on each site per layer.

Definition 3. An *ℓ-layer periodic* random circuit architecture repeats the layout of its layers with period ℓ as the depth increases. Note that gates themselves are independently random at every depth; only the spatial arrangement of the gates is repeated.

Definition 4. A *connected block* of a circuit architecture on N sites is a contiguous sequence of layers such that the gates in the block form a connected graph over all N sites.

These definitions are illustrated in Fig. 1.

Our most general results are for complete architectures and depend on the frequency and size of connected blocks. However, we obtain a more explicit form for the special case of periodic architectures. We obtain each result by reduction to the well-studied 1D brickwork architecture.

Definition 5. The *N-site 1D brickwork* architecture is a complete two-layer periodic random circuit architecture, equipped with an ordering $1, \ldots, N$ of the sites, such that the first layer applies gates on sites $\{2j, 2j+1 \mod N\}$, while the second layer applies gates on sites $\{2j-1 \mod N, 2j\}$. In this paper, the spatial boundary conditions are periodic unless otherwise specified.

E. Main theorems

Theorem 1. Suppose that the N-site 1D brickwork architecture (with either open or periodic boundary conditions) forms an ϵ-approximate t-design after at most

$$k_1 = C(N, q, t) \log \frac{1}{\epsilon} + o_\epsilon\left(\log \frac{1}{\epsilon}\right) \tag{6}$$

periods. Then arbitrary complete ℓ-layer periodic architectures form an ϵ-approximate t-design after at most

$$k_* = \frac{2Nt \log q + \log(1/\epsilon)}{\log(1/s_*)} \tag{7}$$

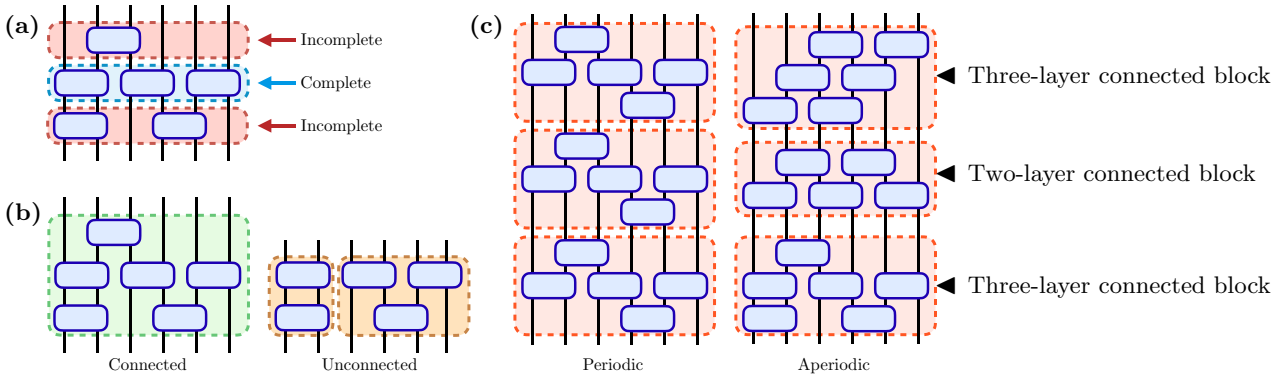

FIG. 1. Different types of layers, blocks, and architectures. (a) The middle layer is complete, since every site is acted on by exactly one gate. The upper and lower layers are both incomplete. (b) The left-hand block is connected, while the right-hand block is made up of two unconnected components. (c) The left-hand architecture is three-layer periodic, containing three repetitions of the same three-layer connected block. The right-hand architecture is aperiodic, consisting of a three-layer connected block, a two-layer connected block, and another three-layer connected block.

periods, where

$$s_* = 1 - \left[1 - \exp\left(-\frac{1}{2C(q,t)}\right)\right]^{\ell-1} \tag{8}$$

and $C(q,t) = \sup_N C(N,q,t)$. We may relax the bound to the more legible

$$k_* = [4\overline{C}(q,t)]^{\ell-1}\left(2Nt\log q + \log\frac{1}{\epsilon}\right) \tag{9}$$

by defining $\overline{C}(q,t) \equiv \max[C(q,t), \frac{1}{2}]$.

Theorem 2. If we do not require that the layers be complete, an ℓ-layer periodic architecture forms an ϵ-approximate t-design after at most

$$k_* = \frac{2Nt\log q + \log(1/\epsilon)}{\log(1/s_*)} \tag{10}$$

periods, where

$$s_* = 1 - \left[1 - \exp\left(-\frac{1}{2C(\sqrt{q},t)}\right)\right]^{\ell-1} \tag{11}$$

and $C(\sqrt{q},t)$ is defined either as in Theorem 1 or, when $\sqrt{q}$ is not an integer, via the generalization described below.

For integer q, $C(q,t)$ can be equivalently defined in terms of the spectrum of the transfer matrix tensor network corresponding to the 1D brickwork architecture (see Figs. 2 and 3 below). In this picture q corresponds to a "coupling constant" in the tensor network. When q is an integer but $\sqrt{q}$ is not, we define $C(\sqrt{q},t)$ analogously by changing the coupling constant to $\sqrt{q}$ in the tensor network. For a formal description, see Appendix C.

In addition, we give an alternative result on incomplete layers in Sec. VII B that depends on $C(q,t)$ instead of $C(\sqrt{q},t)$, but at the cost of multiplying ℓ by a factor $O(\log\log N)$. We also show that the architecture need not actually be periodic.

Theorem 3. The results of Theorems 1 and 2 hold even if the architecture is not periodic, with ℓ replaced by an "average connection depth" defined formally in Theorem 12 below.

We begin with a proof of the periodic, complete case. This proof is simpler and illustrates the essential elements of our strategy. The argument will then be extended to the incomplete and aperiodic cases.

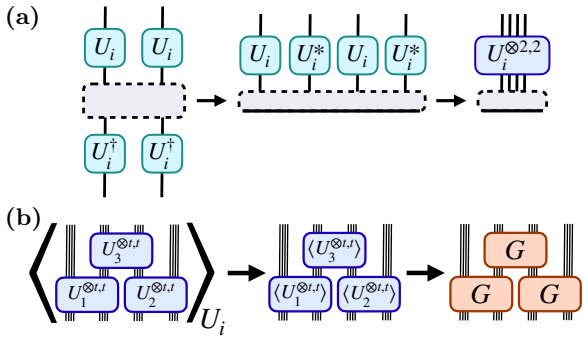

FIG. 2. (a) Folding all t copies of specific gates U_i and $U_i^\dagger$ in channel Φ_{RQC} into a single operator $U_i^{\otimes t,t}$, where $t = 2$. The gray dotted region indicates a density matrix on which the channel acts. This is vectorized in the middle panel, forming a matricization of the $2t$ copies of U_i. (b) The reduction of the t^{th} moment operator $\widehat{\Phi}_{\mathrm{RQC}}$ to a network of identical projection tensors G. Each unitary $U_i^{\otimes t,t}$ can be averaged separately into an independent copy of G.

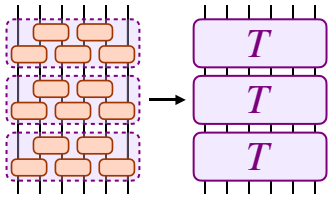

FIG. 3. Breaking up $\widehat{\Phi}_{\mathrm{RQC}}$ for a random quantum circuit (in this case a 1D brickwork architecture) into $k = 3$ copies of the transfer matrix T.

Finally, we show that, conditional on two conjectures, we can omit the dependence on ℓ entirely to obtain a much simpler result.

Theorem 4. Suppose that Conjectures 1 and 2 below hold. Then any architecture that can be divided into

$$k_* = \frac{2Nt\log q + \log(1/\epsilon)}{2\log[(q^2+1)/2q]} \tag{12}$$

connected blocks forms an ϵ-approximate t-design.

F. Known values of $C(q,t)$

Previous works imply the following bounds on $C(q,t)$.

(a) For general parameters, the best known bound is that of Brandão *et al.* [22]. For the case of 1D brickwork circuits [34] where $q \geq 2$, the authors gave

$$C(q,t) = 261\,500\lceil\log_q(4t)\rceil^2 q^2 t^{5+3.1/\log(q)}. \tag{13}$$

For integer $q^2 \geq 2$, the more general form is

$$C(q,t) = 234(q^2+1)e^{\{2.5\log(4)[1+\log(q^2+1)]\}/\log(q)+1} \times \lceil\log_q(4t)\rceil^2 t^{5+5[1+\log(1+q^{-2})]/2\log(q)}. \tag{14}$$

When $q \geq 2$, we can replace the t exponent with $9.5 \geq 5 + \{\log[e(1+q^2)]/\log(q)\}$. Similarly, for $q^2 \geq 2$, we have an exponent of at most 15.2.

(b) For $q = 2$, Eq. (22) of Ref. [23] gives the tighter bound

$$k_1 = \alpha\log^5(t)t^{4+3/\sqrt{\log_2 t}}\left(2Nt + \log_2\frac{1}{\epsilon}\right), \tag{15}$$

where $\alpha = 10^{13}$. This gives

$$C(2,t) = \frac{\alpha}{2\log 2}\log^5(t)t^{4+3/\sqrt{\log_2 t}}. \tag{16}$$

(c) For $t = 2$ and any $q > 1$, Eq. (27) of Ref. [24] gives

$$C(q,2) = \left(2\log\frac{q^2+1}{2q}\right)^{-1}$$

for 1D brickwork circuits with open boundary conditions [35]. We show in Appendix B that periodic boundary conditions improve the bound to

$$C(q,2) = \left(4\log\frac{q^2+1}{2q}\right)^{-1}. \tag{17}$$

(d) In the limit $q \to \infty$, Eq. (36) of Ref. [24] shows that the leading-order term is $C(q,t) = (2\log q/2)^{-1}$ with open boundary conditions. Periodic boundary conditions again tighten [36] this to

$$C(q,t) = \left(4\log\frac{q}{2}\right)^{-1} + o_q(\log^{-1} q). \tag{18}$$

(e) Following the conjecture of Ref. [24], we suspect that the sharp bound is

$$C(q,t) = \left(4\log\frac{q^2+1}{2q}\right)^{-1}. \tag{19}$$

This is analogous to the conjecture of Ref. [24] for the open-boundary case. Numerical evidence is given in Appendix E 1.

II. APPROXIMATE t-DESIGNS AND TENSOR NETWORK PICTURE

The first phase of our proof follows the standard reduction from approximate t-designs to a tensor network of averaged gates [21]. For the sake of completeness and notational clarity, Secs. II and III outline the key steps.

For a random quantum circuit channel Φ_{RQC} formed from a circuit ensemble $U_C \in \varepsilon_C$, the diamond norm difference from the Haar distribution is bounded in terms of the frame potential [24]:

$$\|\Phi_{\mathrm{RQC}} - \Phi_{\mathrm{Haar}}\|_\diamond^2 \leq q^{2Nt}(\mathcal{F}^{(t)}_{\varepsilon_C} - \mathcal{F}^{(t)}_{\mathrm{Haar}}), \tag{20}$$

$$\mathcal{F}^{(t)}_{\varepsilon_C} = \int_{\varepsilon_C^{\otimes 2}} |\mathrm{tr}(U_C^\dagger V_C)|^{2t}\mathrm{d}U_C\mathrm{d}V_C. \tag{21}$$

Both the random quantum circuit channel and the frame potential can be written in terms of the t^{th} moment operator

$$\widehat{\Phi}_{\mathrm{RQC}} = \int_{\varepsilon_C} U_C^{\otimes t,t}\mathrm{d}U_C,$$

where $U_C^{\otimes t,t} \equiv U_C^{\otimes t} \otimes (U_C^*)^{\otimes t}$. This is a matricization of the quantum channel Φ_{RQC}, i.e.,

$$\widehat{\Phi}_{\mathrm{RQC}} \cdot \mathrm{vec}(\rho) = \mathrm{vec}[\Phi_{\mathrm{RQC}}(\rho)]. \tag{22}$$

We also have

$$\begin{aligned}\mathcal{F}_\varepsilon^{(t)} &= \int_{\varepsilon_C^{\otimes 2}} \mathrm{tr}(U_C^{\dagger\otimes t,t} V_C^{\otimes t,t}) \mathrm{d}U_C \mathrm{d}V_C \\ &= \mathrm{tr}(\widehat{\Phi}_{\mathrm{RQC}}^\dagger \widehat{\Phi}_{\mathrm{RQC}}).\end{aligned} \tag{23}$$

We assume that U_C consists of two-site unitary gates U_i drawn independently from the Haar distribution over $U(q^2)$. Since distinct gates are independent, we can average over each gate separately (Fig. 2). The averaging joins the $2t$ copies of each gate that appear in $U_C^{\otimes t,t}$ into a single operator G. The action of G depends only on t and the number of sites on which U_i acts. Operator $\widehat{\Phi}_{\mathrm{RQC}}$ becomes the contraction of a tensor network in the shape of the original circuit U_C, but with each $U_i^{\otimes t,t}$ replaced with its average G.

The individual G's can be written in terms of single-site *permutation states*. Given a permutation $\sigma \in S_t$, we define a particular maximally entangled state on t pairs of sites

$$|\sigma\rangle = q^{-t/2} \sum_{\vec{i} \in \mathbb{Z}_q^t} |\vec{i}\rangle |\sigma(\vec{i})\rangle. \tag{24}$$

We call a tensor power of a permutation state on m sites $|\sigma\rangle^{\otimes m} = |\sigma\rangle|\sigma\rangle \cdots |\sigma\rangle$ a *uniform permutation state*.

Theorem 5. Let U be an m-site Haar-random unitary. Let G be the expected value of the corresponding moment operator $U_i^{\otimes t,t}$. Then G is a projector onto the space spanned by the uniform permutation states $|\sigma\rangle^{\otimes m}$.

A proof may be found in Ref. [33]. In particular, we see that the t^{th} moment operator for the Haar distribution over all N sites is just the orthogonal projector on to the globally uniform permutation states on all N sites. Furthermore, these states also span the unit eigenspace of the moment operator of any architecture.

Lemma 1. If the circuit architecture is connected, the unit eigenspace of $\widehat{\Phi}_{\mathrm{RQC}}$ is spanned by the globally uniform permutation states $|\sigma\rangle^{\otimes N}$.

Proof. The support of the distribution over the unitaries induced by a random architecture is a universal gate set if and only if the architecture is connected [37,38]. Suppose that we apply k repetitions of the circuit architecture. As $k \to \infty$, Emerson *et al.* [19] showed that the induced measure on $U(q^N)$ converges to the Haar measure. It follows that the corresponding moment operator $\widehat{\Phi}_{\mathrm{RQC}}^k$ converges to that of the Haar measure, which is the projector onto the span of the uniform permutation states. But, since $\widehat{\Phi}_{\mathrm{RQC}}$ is norm nonincreasing, $\lim_{k\to\infty} \widehat{\Phi}_{\mathrm{RQC}}^k$ is a projector on to the unit eigenspace of $\widehat{\Phi}_{\mathrm{RQC}}$, so the unit eigenspace of $\widehat{\Phi}_{\mathrm{RQC}}$ must be the same as that of $\widehat{\Phi}_{\mathrm{Haar}}$. Theorem 5 completes the argument. ■

III. TRANSFER MATRIX AND THE SPECTRAL GAP

We now specialize to the case of ℓ-layer periodic architectures, again following the standard techniques found in Ref. [21]. In this case we can define a transfer matrix T by contracting together the projectors G_i, $i \in \{1, \ldots, \ell N/2\}$, of the moment operator $\widehat{\Phi}_{\mathrm{RQC}}$ corresponding to a single period of the architecture, as shown in Fig. 3. We show that the approximate t-design time is controlled by the singular value spectrum of T.

If there are k periods of the architecture, the moment operator corresponds to the kth power of the transfer matrix, so by Eq. (23), the frame potential is

$$\mathcal{F}_\varepsilon = \mathrm{tr}(T^{\dagger k} T^k) = \|T^k\|_F^2. \tag{25}$$

Theorem 6. Consider a connected periodic architecture on N sites. After k periods, the frame potential is at most

$$\mathcal{F}_\varepsilon \le \mathcal{F}_{\mathrm{Haar}} + q^{2Nt} s_*^{2k}, \tag{26}$$

where s_* is the largest nonunit singular value of the transfer matrix T.

Proof. Starting from

$$\mathcal{F}_\varepsilon = \|T^k\|_F^2,$$

we can use Theorem 17, proven in Appendix A, to see that

$$\|T^k\|_F^2 \le m_1 + (d^2 - m_1 - m_0) s_*^{2k}, \tag{27}$$

where d^2 is the dimension of T, m_1 is the dimension of its unit eigenspace, and m_0 is the dimension of its zero eigenspace. As long as the circuit is connected, by Lemma 1, it shares the same unit eigenstates as the Haar distribution, so $m_1 = \mathcal{F}_{\mathrm{Haar}}^{(q^N)}$. The first layer consists of two-site gates, each of which has only $\mathcal{F}_{\mathrm{Haar}}^{(q^2)}$ nonzero eigenvalues. A lower bound on m_0 is thus $d^2 - (\mathcal{F}_{\mathrm{Haar}}^{(q^2)})^N$. The total dimension for T is $d^2 = q^{2Nt}$, so we find that

$$\mathcal{F}_\varepsilon \le \mathcal{F}_{\mathrm{Haar}}^{(q^N)} + [(\mathcal{F}_{\mathrm{Haar}}^{(q^2)})^N - \mathcal{F}_{\mathrm{Haar}}^{(q^N)}] s_*^{2k}. \tag{28}$$

Of course, $\mathcal{F}_{\mathrm{Haar}}$ is just the dimension of the unit eigenspace, which is at most $t!$ (since it is spanned by the uniform permutation states) and also at most q^{2Nt} (since that is the total dimension of T). In addition, it is known to

be equal [24] to $t!$ for $t \leq d$. From Eq. (21), it is clear that the frame potential is a monotonically increasing function of t. Together these imply that

$$\min(t!, d!) \leq \mathcal{F}^{(d)}_{\text{Haar}} \leq \min(t!, d^{2t}), \tag{29}$$

so that

$$(\mathcal{F}^{(q^2)}_{\text{Haar}})^N - \mathcal{F}^{(q^N)}_{\text{Haar}} \leq \min(t!, q^{2t})^N - \min[t!, (q^N)!].$$

The remaining algebra will be easier if we further relax this to

$$(\mathcal{F}^{(q^2)}_{\text{Haar}})^N - \mathcal{F}^{(q^N)}_{\text{Haar}} \leq q^{2Nt}. \tag{30}$$

Without this relaxation, all of our bounds change by

$$2Nt \log q \to \log[(\mathcal{F}^{(q^2)}_{\text{Haar}})^N - \mathcal{F}^{(q^N)}_{\text{Haar}}].$$

This latter bound is tighter for small t, but the former bound is more convenient to interpret, so we prefer the simplified form. Equation (28) then becomes

$$\mathcal{F}_\varepsilon \leq \mathcal{F}_{\text{Haar}} + q^{2Nt} s_*^{2k}. \tag{31}$$

This completes the proof. ■

Applying Theorem 6 to Eq. (20) gives a bound on the rate at which a random circuit architecture approaches a t-design in terms of the subleading singular value (SSV) of the transfer matrix:

$$\|\Phi_{\text{RQC}} - \Phi_{\text{Haar}}\|_\diamond \leq q^{2Nt} s_*^k. \tag{32}$$

It follows that the number of periods required to push the diamond norm error below ϵ can be upper bounded in terms of s_* by

$$k_* = \frac{2Nt \log q + \log(1/\epsilon)}{\log(1/s_*)}, \tag{33}$$

which is already part of Theorem 1. It remains only to bound s_*, i.e., the spectral gap of T.

IV. BOUNDING THE SPECTRAL GAP

In order to apply Theorem 6, we must bound the largest nonunit singular value T. We first show that this s_* is related to the geometry of the unit eigenspaces of each layer. Later, we study the relationship between this geometry and the architecture of the circuit to derive our final bound.

We can interpret T as a product of orthogonal projection operators. Because each projector is norm nonincreasing, any subspace whose norm is decreased by some large amount by the first few projectors cannot contain a large singular value. If $P_i \cdots P_1$ does not have a large singular value then any large singular value of $P_{i+1} \cdots P_1$ must arise from vectors that were nearly unit eigenvectors of $P_i \cdots P_1$, while still being orthogonal to the unit eigenvectors of $P_{i+1} \cdots P_1$. This allows us to construct an inductive bound based on the relative geometry of the subspaces to which the P_i project.

Theorem 7. Consider some set of subspaces X_i, $i \in \{1, \ldots, n\}$, of a Hilbert space. Let P_i be the orthogonal projector on to X_i and Q_i the orthogonal projector on to $\bigcap_{j=1}^{i} X_j$. Define $T = P_n \cdots P_2 P_1$, and let s_* be the largest nonunit singular value of T_n. Then we have the bound

$$s_*^2 \leq 1 - \prod_{i=2}^{n} (1 - \mathfrak{s}_i^2), \tag{34}$$

where $\mathfrak{s}_i$ is the largest nonunit singular value of $P_i Q_{i-1}$.

The proof of this theorem is given in Appendix A. In our case, the P_i will be the layers of the transfer matrix. We call the $\mathfrak{s}_i$ the *layer-restricted singular values*. Our first goal is to characterize the Q_i, which we term the *intermediate unit eigenspace projectors*.

A. Cluster-merging picture

To bound the layer-restricted singular values, we must first understand the intermediate unit eigenspace to which Q_i projects. From Theorem 5, we know that the unit eigenstates of a single gate are the uniform permutation states $|\sigma\rangle^{\otimes m}$. Moreover, each gate is norm nonincreasing, so a state whose norm is reduced by any individual gate must not be a unit eigenstate. This leads to the following result.

Lemma 2. Let G_i, $i \in \{1, \ldots, m\}$, be some gates within the transfer matrix that form a connected network on some n sites. The unit eigenspace of the contracted network $M = \prod_{i=1}^{m} G_i$ is spanned by the uniform permutation states $\{|\sigma\rangle^{\otimes n}, \sigma \in S_t\}$.

Proof. We may interpret M as the transfer matrix of a (possibly incomplete) random quantum circuit architecture with the same layout as the G_i. Since the G_i are connected, the corresponding quantum circuits are also connected. We may then apply Lemma 1. ■

This leads to the *cluster-merging picture* (Fig. 4). Take some sequence of complete layers $P_1 \cdots P_l$.

These layers may connect all the sites, or they may connect only certain subsets of the sites. Call each subset of sites that is connected by $P_1 \cdots P_l$ a *cluster*. Recall that the intermediate unit eigenspace projector Q_l is defined to be the orthogonal projector on to the unit eigenspace of the product $P_l \cdots P_1$.

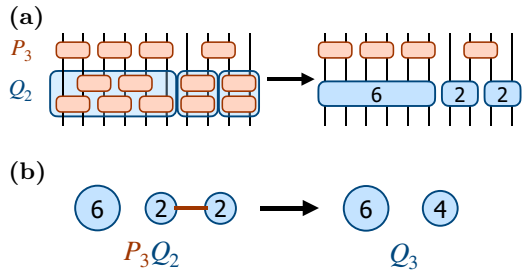

FIG. 4. The cluster-merging picture for calculating the third layer-restricted subleading singular value ϑ_3, i.e., the subleading singular value of P_3Q_2. (a) The first two layers are collected into clusters (blue) of sites that are connected by the first two layers of gates. The intermediate unit eigenspace projector Q_2 is then a tensor product of projectors on each cluster. (b) (*Left*) Graphical representation of P_3Q_2 in the cluster-merging picture. Each cluster in Q_2 becomes a node labeled by the number of sites it contains. Gates in P_3 that join distinct clusters become edges. Gates in P_3 that join two sites within the same cluster are ignored. (*Right*) Each connected component of this graph is merged into a single cluster of the next intermediate unit eigenspace projector Q_3.

By Lemma 2, the space to which Q_l projects is the tensor product of the spans of the uniform permutation states on each cluster. As we include more layers by increasing l, the clusters merge and the unit eigenspace to which Q_l projects shrinks. Eventually, the whole circuit has been connected into a single cluster, at which time the unit eigenspace is the span of the globally uniform permutation states $|\sigma\rangle^{\otimes N}$.

Our goal is to compute ϑ_l, which is the largest nonunit singular value of P_lQ_{l-1}. This can be accomplished by constructing the *cluster-merging graph* that uniquely determines ϑ_l. To build the graph, let each cluster of sites connected by $P_{l-1}\cdots P_1$ be a node, with weight equal to the number of sites in the cluster. Each gate of P_l that joins two distinct clusters is mapped to an edge joining the corresponding nodes. Gates of P_l that join two sites within the same cluster do not influence ϑ_l and can be ignored. The layer-restricted singular values ϑ_l depend only on the graph topology and the node weights, not on any other details of the architecture.

V. REDUCTION OF EACH LAYER TO 1D BRICKWORK LOOPS

Our next goal is to obtain an architecture-independent upper bound on the ϑ_i. We begin by identifying a set of rules for rearranging a cluster-merging graph into a standardized form without decreasing ϑ_i.

1. Structure of the graph

Lemma 3. For a complete ℓ-layer periodic architecture, the cluster-merging graph for each layer above the first has nodes of even weight and even degree.

Proof. We first show that the clusters are of even size. The first layer creates clusters of size 2. Subsequent layers create clusters by merging these size-2 clusters together, so later cluster sizes are also even.

We next show that each cluster has an even number of external connections. Suppose that a cluster of size m has n_i internal gates and n_e external gates. Since each site is acted on by exactly one two-site gate, we have $2n_i + n_e = m$, and so n_e is even. The nodes are thus of even degree. ■

2. Cluster-merging bound

Our goal is to bound ϑ_i for generic cluster-merging graphs in terms of cluster-merging graphs of some standard form. First consider the unit eigenspace.

Lemma 4. The unit eigenspace of a connected cluster-merging graph is spanned by the globally uniform permutation states.

This is just a special case of Lemma 2. We can now show a lower bound.

Lemma 5. Let A,B be two nodes of a connected cluster-merging graph. If we merge A and B into a single node AB, the subleading singular value of the graph does not increase.

Proof. Let ϑ, ϑ' be the subleading singular values of the old and new graphs, respectively. Let X, X' be the subspaces projected to by the nodes of the old and new graphs. Let P be the projector corresponding to the edges, which are the same for both graphs.

By Lemma 4, the unit eigenspaces of PQ and PQ' are the same. Call this subspace Z. We may write

$$\vartheta = \max_{\{v\in X | v\perp Z\}} \frac{\|Pv\|}{\|v\|} \tag{35}$$

and

$$\vartheta' = \max_{\{v\in X' | v\perp Z\}} \frac{\|Pv\|}{\|v\|}. \tag{36}$$

Subspace X is spanned by the cluster-uniform states over the old nodes, while X' is spanned by the cluster-uniform states over the new nodes. A cluster-uniform permutation state on the new node AB is of the form $|\sigma\rangle^{\otimes(|A|+|B|)}$, which is also cluster uniform on A and B individually. In other words, $X' \subseteq X$. Since ϑ' is the maximum of the same function over a smaller space, we obtain $\vartheta' \leq \vartheta$. ■

If we run our lower bound in reverse, we are led to the following two rewriting rules that give upper bounds on subleading singular values of cluster-merging graphs.

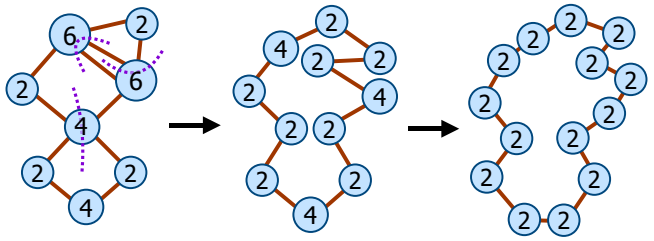

FIG. 5. Algorithm for bounding a cluster-merging graph by a 1D brickwork structure. We first use Lemma 6 to split clusters (without joining the halves with an edge) along a Eulerian circuit in our graph. We then use Lemma 7 to break up larger clusters (joining the halves with an edge) until all clusters in the loop have size 2. Each step in this process does not decrease the subleading singular value of the graph.

Lemma 6. If we split a node and the two sides remain part of the same connected graph, the subleading singular value does not decrease.

Lemma 7. If we split a node and add a new link to keep the graph connected, the subleading singular value does not decrease.

Proofs of Lemmas 6 and 7. These both correspond to applying Lemma 5 backwards. ■

Note that the number of links connected to a cluster cannot exceed the number of sites it contains, so we can only apply the latter lemma to clusters with at least two unoccupied sites. We are now ready to prove the main result of this section.

Theorem 8. Consider any cluster-merging graph on n sites with SSV $\mathfrak{z}$. Let $s_{1D}(m)$ be the SSV for a 1D brickwork loop on m sites. We have

$$\mathfrak{z} \leq \max_{m\leq n} s_{1D}(m). \tag{37}$$

Proof. Our goal is to apply the graph rewriting rules repeatedly to upper bound a cluster in a standard form. First suppose that the graph is connected. Since nodes are of even degree, there exists a Eulerian circuit through the cluster-merging graph. We can split nodes along this Eulerian circuit using Lemma 6 until our graph structure is a single loop, again with nodes of even degree (Fig. 5, first process). We can then apply Lemma 7 repeatedly within each node to split each node into many nodes, each of size 2 (Fig. 5, second process). When we are done, we have exactly a 1D brickwork loop on m sites. These transformations cannot decrease the SSV, so the brickwork SSV is an upper bound on the original SSV.

Now consider any disconnected graph. We can apply the argument above to bound each connected component by a loop graph. The corresponding operator is a tensor product of loop operators, and the largest singular value of each loop operator is 1. The subleading singular value of the whole operator is then just the largest subleading singular value of any of the connected components. Each connected component is a brickwork loop of size at most $m \leq n$. ■

VI. SPECTRAL GAP OF 1D BRICKWORK LOOPS

In order to extract a useful result from Theorem 8, we need a bound on the spectral gap of the 1D brickwork architecture. In Sec. V we showed that the t-design depth is controlled by the spectral gap of the transfer matrix. We now reverse that argument in order to obtain a bound on the spectral gap in terms of the t-design depth. This will allow us to convert any result on 1D brickwork t-design depths into bounds on the spectral gap.

Theorem 9. Suppose that the 1D brickwork architecture on N sites (with either open or periodic boundary conditions) forms an ϵ-approximate t-design after

$$k_1 = C(N,q,t)\log\frac{1}{\epsilon} + o_\epsilon\left(\log\frac{1}{\epsilon}\right) \tag{38}$$

periods (corresponding to depth $2k_1$) for some function $C(N,q,t)$. Then the largest nonunit singular value s_{1D} of the corresponding transfer matrix T_t is bounded by

$$s_{1D} \leq e^{-1/2C(q,t)}, \tag{39}$$

where $C(q,t) = \sup_N C(N,q,t)$.

Proof. Define

$$\Delta(k,t) = \Phi_{\mathrm{RQC},k,t} - \Phi_{\mathrm{Haar}}.$$

From the definition of the diamond norm,

$$\|\Delta\|_\diamond \geq \|\Delta\|_{1\to 1} \geq \frac{\|\Delta(\rho)\|_1}{\|\rho\|_1} \tag{40}$$

for any operators ρ. Let λ be the largest eigenvalue of the Δ and choose ρ to be the corresponding eigenvector, so that $\Delta(\rho) = \lambda\rho$. Then we obtain

$$\|\Delta\|_\diamond \geq |\lambda|. \tag{41}$$

Note that the leading eigenvalue of Δ is exactly the subleading eigenvalue of Φ_{RQC}, since the unit eigenspace of Φ_{RQC} is exactly canceled by Φ_{Haar}. Furthermore, $\hat{\Phi}_{\mathrm{RQC}} =$

T_t^k, so we have

$$\lambda = \lambda_*^k, \tag{42}$$

where λ_* is the subleading eigenvalue of T. If we choose $k = k_1$ so that

$$\|\Delta(k_1, t)\|_\diamond < \epsilon,$$

we obtain

$$|\lambda_*| < \epsilon^{1/k_1}. \tag{43}$$

Inverting Eq. (38) gives

$$\log \frac{1}{\epsilon} = \frac{k_1}{C(N, q, t)} + o_{k_*}(k_*) \tag{44}$$

or

$$\epsilon = e^{-k_1/C(N,q,t) + o_{k_1}(k_1)}. \tag{45}$$

We can now take the limit of small ϵ or equivalently large k_1 to obtain

$$\epsilon = (e^{-1/C(N,q,t)})^{k_1}, \tag{46}$$

which implies that $|\lambda_*| \leq e^{-1/C(q,t)}$. Theorem 14 in Appendix A tells us that $s_{1D} = \sqrt{\lambda_*}$, so we find that $s_{1D} \leq e^{-[1/2C(q,t)]}$. ■

Theorem 10. Let $s_{1D,\text{open}}(N)$ be the subleading singular value of the N-site 1D brickwork architecture with open boundary conditions, and let $s_{1D,\text{periodic}}(N)$ be the same for the periodic-boundary-condition case. Then $s_{1D,\text{periodic}}(N) \leq s_{1D,\text{open}}(N)$.

This result allows us to also use bounds derived for open brickwork architectures in Theorem 8. It follows directly from the following more general rewriting rule for cluster-merging graphs.

Lemma 8. Consider a connected cluster-merging graph. If we add a new edge to the graph, the subleading singular value $\mathfrak{s}$ does not increase.

Proof. Let P and Q be the edges and nodes of the original graph, and let R be the new link. Let X be the unit eigenspace of Q and Z be the unit eigenspace of PQ. Since the graph was already connected without R, the unit eigenspace of R includes Z, so Z is also the unit eigenspace of the new graph RPQ. The original singular value is

$$\mathfrak{s} = \max_{\{v \in X | v \perp Z\}} \frac{\|Pv\|}{\|v\|}, \tag{47}$$

while the new singular value is

$$\mathfrak{s}' = \max_{\{v \in X | v \perp Z\}} \frac{\|RPv\|}{\|v\|}. \tag{48}$$

Since R is an orthogonal projector, $\|RPv\| \leq \|Pv\|$, so $\mathfrak{s}' \leq \mathfrak{s}$. ■

Given an open-boundary-condition transfer matrix, we can add a link to obtain the transfer matrix of a periodic-boundary-condition brickwork on the same number of sites. Adding a link can only decrease $\mathfrak{s}$, which completes the proof of Theorem 10.

VII. APPROXIMATE t-DESIGN DEPTHS

We are now ready to prove Theorem 1, which addresses the case of complete periodic architectures. We then proceed to extend our results to incomplete and aperiodic architectures. Figure 1 illustrates the relevant categories.

A. Complete periodic architectures

Suppose that the N-site 1D brickwork architecture forms an ϵ-approximate t-design in depth at most

$$2C(N, q, t) \log \frac{1}{\epsilon} + o_\epsilon\left(\log \frac{1}{\epsilon}\right). \tag{49}$$

Define $C(q, t) = \sup_N C(N, q, t)$. By Theorem 9, the subleading singular value of T_t for a brickwork loop is upper bounded by

$$s_{1D} \leq e^{-1/2C(q,t)}. \tag{50}$$

From Theorem 8, it follows that the SSV for any cluster-merging graph is bounded by the same value. Substituting into Theorem 7, we see that the subleading singular value for a complete ℓ-layer block is bounded by

$$s_* \leq 1 - (1 - e^{-1/2C(q,t)})^{\ell - 1}. \tag{51}$$

We can substitute into Eq. (33) to see that the critical period count is then bounded by

$$k_* = \frac{2Nt \log 2q + \log(1/\epsilon)}{\log\{[1 - (1 - e^{-1/2C(q,t)})^{\ell-1}]^{-1}\}}. \tag{52}$$

This completes the proof of Theorem 1.

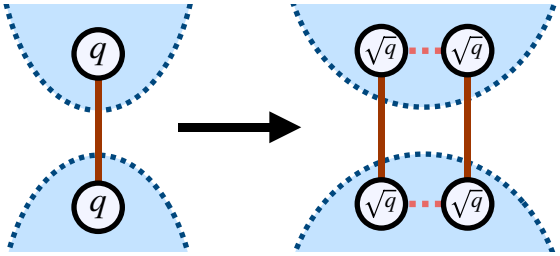

FIG. 6. A two-site gate being split under the $X_q \to X_{\sqrt{q}}^{\otimes 2}$ isomorphism. Each site on X_q becomes a pair of twinned sites (pink dotted lines) on $X_{\sqrt{q}}$. Note that the dimensionality of each site does not necessarily decrease—only the metric between the basis vectors changes. If we assume that the two-site gate spans two different clusters (blue dotted circles) then the layer-restricted SSV of a hyperedge connecting all four split sites is equal to two pairs of edges connecting nontwinned sites.

B. Incomplete layers

For complete architectures, we have

$$s_*^2 \leq 1 - [1 - s_{1D}(q)]^{\ell - 1} \tag{53}$$

by reducing each cluster-merging graph to a Eulerian cycle. If the circuit is incomplete, the Eulerian cycle is not guaranteed to exist. Instead, an incomplete circuit with local dimension q has subleading singular value bounded by that of a complete circuit with local dimension $\sqrt{q}$.

We may, without loss of generality, insert single-site gates so that each site is acted upon by exactly one gate per layer. These single-site gates can be absorbed into a two-site gate either above or below without changing the Haar measure over that two-site gate. This takes an incomplete circuit to a "complete" circuit with a mixture of one-site and two-site gates.

Now suppose that $\sqrt{q}$ is an integer. We can then split each site of dimension q into a pair of sites of dimension $\sqrt{q}$. This gives us a complete circuit containing a mixture of two- and four-site gates. Suppose that we split each four-site gate lengthwise into a pair of two-site gates. It seems intuitive that this operation should decrease the rate of scrambling, not increase it. And the resulting circuit is a complete circuit of two-site gates with local dimension $\sqrt{q}$, so its cluster-merging graphs have Eulerian cycles. In Appendix C, we formalize this intuition and show that the same idea can be applied even in the case where $\sqrt{q}$ is not an integer. The site-splitting process is illustrated in Fig. 6 in Appendix C.

This allows us to reduce a bound on incomplete layers to a 1D brickwork bound on $\sqrt{q}$; specifically,

$$s_*^2 \leq 1 - [1 - s_{1D}(\sqrt{q})]^{\ell - 1} \tag{54}$$

for incomplete architectures.

We may use any proof of $C(q,t)$ if $\sqrt{q}$ is an integer. For noninteger $\sqrt{q}$, however, the definition of s_{1D} is trickier. It is not necessarily true that generic bounds on $C(q,t)$ that have been derived for integer $q \geq 2$ can be analytically continued to $C(\sqrt{q},t)$. Many of the proofs in the literature can easily be extended to noninteger $q > 1$ (see Sec. I F). However, not all of these proofs can be easily extended, e.g., the proof of Ref. [23] applies only to $q = 2$. The bound of Ref. [22] can be extended, but the scaling goes from about t^{10} to about t^{16}. And there may be future improved strategies for bounding 1D brickwork that work only for integer q. We thus also give the following bound, which gives us a reduction in terms of $C(q,t)$ instead of $C(\sqrt{q},t)$.

Theorem 11. For an architecture with blocks of ℓ layers,

$$k_* = [4\overline{C}(q,t)]^{x(N)\ell - 1}\left(2Nt \log q + \log \frac{1}{\epsilon}\right), \tag{55}$$

where the expansion coefficient

$$x(N) = 8\lceil \log_2 \lfloor \log_2 (N+1) \rfloor \rceil + 2. \tag{56}$$

The proof is given in Appendix D.

C. Aperiodic architectures

So far, we have restricted ourselves to periodic architectures to simplify the exposition. However, our results generalize quite directly to aperiodic architectures.

Theorem 12. Define $C(q,t)$ as in Theorem 1. For a not-necessarily-periodic L-layer circuit U_C, choose a decomposition

$$U_C = V_{k+1} U_k V_k U_{k-1} V_{k-1} \cdots U_1 V_1, \tag{57}$$

where each U_i is an ℓ_i-layer connected block and each V_i is some contiguous block of layers that may not connect all the sites. The architecture is an ϵ-approximate t-design if

$$k \geq \frac{2Nt \log q + \log(1/\epsilon)}{\log(1/s_*)}, \tag{58}$$

where we have defined the effective averaged singular value

$$s_* = 1 - (1 - e^{-1/2C(\sqrt{q},t)})^{\bar{\ell}-1}$$

in terms of the mean block size

$$\bar{\ell} = \frac{1}{k}\sum_{i=1}^{k} \ell_i.$$

If we also require that the layers be complete, we can replace $\sqrt{q}$ with q. Note that k counts the total number

of times the circuit is connected, while $\bar{\ell}$ is the size of the typical connected block. We see that the (ϵ, t)-design depth is controlled by the frequency and size of the connected blocks. However, for an aperiodic architecture, $\bar{\ell}$ might depend on depth, so this does not give us any explicit expression for the critical depth.

Proof of Theorem 12. A decomposition of the circuit into blocks of layers induces a decomposition of the associated t-fold channel as

$$\hat{\Phi}_{\mathrm{RQC}} = R_k T_k R_{k-1} T_{k-1} \cdots R_1 T_1,$$

where the R_i are some set of norm-nonincreasing transfer matrices. Let $s_*^{(i)}$ be the subleading singular value of each T_i. Equation (51) still applies for each $s_*^{(i)}$, with ℓ replaced by ℓ_i [39]. Theorem 16 in Appendix A shows that the R_i are essentially irrelevant. Equation (32) now becomes

$$\|\Phi_{\mathrm{RQC}} - \Phi_{\mathrm{Haar}}\|_\diamond \leq q^{2Nt} \prod_{i=1}^{k} s_*^{(i)}. \tag{59}$$

The condition to obtain an ϵ-approximate t-design is then

$$\epsilon \leq q^{2Nt} \prod_{i=1}^{k} [1 - (1 - s_{1\mathrm{D}})^{\ell_i - 1}]. \tag{60}$$

We may rearrange this as

$$2Nt \log q + \log \frac{1}{\epsilon} \geq - \sum_{i=1}^{k} \log[1 - (1 - s_{1\mathrm{D}})^{\ell_i - 1}]. \tag{61}$$

Furthermore, $-\log(1 - c^x)$ is convex, so, by Jensen's inequality,

$$2Nt \log q + \log \frac{1}{\epsilon} \geq -k \log[1 - (1 - s_{1\mathrm{D}})^{\bar{\ell} - 1}]. \tag{62}$$

This completes the proof. ■

This formula implies Theorem 2 as a special case. Furthermore, it is also simple for *regularly connected* architectures for which all of the ℓ_i are equal.

There is still a question of the choice of decomposition in Eq. (57). This decomposition is not unique; different choices of decomposition will give different bounds. Note in particular that V_i may be empty, which corresponds to the identity.

Furthermore, the optimal decomposition may depend on q and t. For example, consider an architecture consisting of alternating two-layer and four-layer connected blocks. If we count all the blocks then $\bar{\ell} = 3$ and the depth is $d = 3k$. But if we count only the two-layer blocks and lump the four-layer blocks into the V_i, we obtain $\bar{\ell} = 2$ and $d = 6k$. The former is better for small C and the latter for large C, with a crossover point $C(\sqrt{q}, t) \approx 1.157$.

VIII. FURTHER EXTENSIONS

A. ℓ-independent bound

It is interesting to note that our Theorem 1 gives a bound that loosens as the period ℓ increases. It seems likely that the approximate t-design depth actually decreases [40] with ℓ for certain well-connected structures, such as higher-dimensional brickwork architectures. But the scaling of our bound, which is determined by *worst-case architectures*, suggests that there may exist strange "tenuously connected" architectures at larger ℓ.

It seems intuitively clear that the open-boundary-condition brickwork architecture is in some sense the most "spread out" arrangement of gates possible. Any other connected architecture must involve more total gates, and graph distances between sites must be shorter. Indeed, optimization of the subleading singular value of small circuits using simulated annealing (see Appendix E 2) have failed to find any of these "tenuously connected" architectures: all subleading singular values are bounded by that of the open-boundary-condition brickwork. This motivates the following conjecture.

Conjecture 1. Every connected architecture on N sites has subleading singular value $s_* \leq s_{1\mathrm{D,open}}(N, q, t)$.

The immediate consequence of this conjecture is that there is a universal t-design *connection count* that does not depend on the circuit architecture even via ℓ. In particular, any circuit architecture that can be divided into

$$k_* = \frac{2Nt \log q + \log(1/\epsilon)}{\log(1/s_{1\mathrm{D,open}})} \tag{63}$$

connected blocks forms an (ϵ, t)-design.

To obtain Theorem 4, we also follow Ref. [24] in making the following guess.

Conjecture 2. The subleading singular value of the open-boundary-condition 1D brickwork architecture is

$$s_{1\mathrm{D,open}}(q, t) = \frac{2q}{q^2 + 1}. \tag{64}$$

Numerical evidence for this formula is given in Appendix E 1.

B. Nondeterministic architectures

Our theorems focus on deterministic architectures, in which the contents of the gates are random, but their arrangement is fixed. Another interesting class of ensembles is nondeterministic architectures, in which the locations of the gates are also random [22]. In this case, the t-design property is obtained by averaging over both the spatial arrangement of the gates and their content. Bounds

on nondeterministic and deterministic architectures can often be related to each other by the union bound or detectability lemma [23,41].

We can also show more directly that bounds for particular spatial structures imply bounds for averages over ensembles of structures. We use the triangle inequality on $\|\Phi_{\mathrm{RQC}} - \Phi_{\mathrm{Haar}}\|_\diamond$, where Φ_{RQC} is drawn from some distribution ρ_Φ. Then

$$\|\langle\Phi_{\mathrm{RQC}}\rangle_{\rho_\Phi} - \Phi_{\mathrm{Haar}}\|_\diamond \leq \langle\|\Phi_{\mathrm{RQC}} - \Phi_{\mathrm{Haar}}\|_\diamond\rangle_{\rho_\Phi}. \tag{65}$$

In particular, we can apply Eq. (60) with the right-hand side replaced with its average over the architecture. We can use Jensen's inequality again to see that

$$\log\langle\|\Phi_{\mathrm{RQC}} - \Phi_{\mathrm{Haar}}\|_\diamond\rangle_{\rho_\Phi} \leq \langle\log\|\Phi_{\mathrm{RQC}} - \Phi_{\mathrm{Haar}}\|_\diamond\rangle_{\rho_\Phi}, \tag{66}$$

so we can also apply Eq. (62) with the right-hand side replaced by its average.

The connection count k and mean block size $\bar{\ell}$ will both differ between realizations, and it is not clear that there is any general strategy for calculating their distribution. Presumably, they are not independent, so we cannot solve Eq. (62) for $\langle k\rangle$. An interesting open question is whether there exists any simple relationship between the averaged bound and the distribution from which the circuit structure is drawn.

A commonly studied example is that of gates that are drawn sequentially from some uniform distribution over some set of pairs of sites [22,23,28]. In the worst case, such circuits may require $N-1$ nontrivial layers to connect all the sites [42]. However, *typical* instances probably require far fewer distinct layers, perhaps as few as $O(1)$. Further work may wish to explore the relationship between bounds obtained by this strategy and known bounds for such architectures [28].

If we assume that Conjecture 1 holds, the nondeterministic case becomes more tractable. We find that it forms an approximate t-design when

$$\langle k\rangle \geq \frac{2Nt\log q + \log(1/\epsilon)}{\log(1/s_{\mathrm{1D,open}})}. \tag{67}$$

Suppose that we have n_g gates sampled independently and identically from the uniform distribution over the edges of some large connected graph. There is quite a bit more we can say about the relationship between n_g and $\langle k\rangle$. The fully connected graph, for example, undergoes a percolation phase transition and becomes connected after $O(N\log N)$ gates, so $n_g = O(N\log N)\langle k\rangle$ [43]. This corresponds to an approximate t-design threshold size of $n_g = O(N^2\log N)$. For $t=2$, this is worse than the known bound by a factor of $\log N$ [20]. From the coupon collector problem, we see that the linear graph also becomes connected after $O(N\log N)$ gates, so its threshold is the same [44]. Again, a bound that is better by a factor of $\log N$ (and that does not rely on Conjecture 1) is already known [22].

For a graph with E edges, we can again use the coupon collector problem to see that it must be connected after $O(E\log E)$ gates. Since $E \leq O(N^2)$, this suggests that every graph gives an approximate t-design threshold size of at most $n_g = O(N^3\log N)$. For graphs that admit Hamiltonian paths, this result is also weaker by $\log N$ than what was already known, but for graphs without Hamiltonian paths, it would be a new result [28].

C. Highly connected architectures

Previous work on random circuits [22–24,40] has focused largely on brickwork architectures. Brickwork architectures are in some sense exceptionally well connected, so they should be expected to converge to the Haar distribution relatively quickly. Indeed, Harrow and Mehraban [25] showed that certain higher-dimensional brickwork circuits approach the Haar measure at a rate that increases with geometric dimension, which corresponds to increasing the period ℓ of the architecture. Here we suggest an extension of our techniques to obtain tighter bounds for such special architectures.

The first observation is that highly connected cluster-merging graphs give small ϑ_i. Second, certain families of N-site architectures form clusters of sizes that scale with N after only a few layers. If we join two m-site clusters with m gates, in Appendix F we show that the subleading value is $q^{-\Omega(m)}$ as $m\to\infty$. In other words, any layers that only join clusters with an extensive number of edges do not contribute to our bound.

For the specific case of higher-dimensional brickwork circuits and $t=2$, we can give a more explicit calculation. Consider a D-dimensional lattice of $N = L^D$ points at positions $\vec{x}\in\{1,\ldots,L\}^D$, with $\hat{e}_i$ being the unit vector in direction i, and connect the lattice with the following sequence of layers.

(a) At layers $1\leq i\leq D$, we apply gates on pairs $(\vec{x},\vec{x}+\hat{e}_i)$ for each $\vec{x}$ such that x_i is even (all addition of components is performed modulo L, to make this architecture periodic).
(b) At layers $D+1\leq i\leq 2D$, we apply gates on pairs $(\vec{x},\vec{x}+\hat{e}_{i-D})$ for each $\vec{x}$ such that x_{i-D} is odd.

This is a higher-dimensional generalization of the brickwork architecture that emphasizes accessing all dimensions as quickly as possible, instead of repeating the one-dimensional brickwork across multiple directions.

The first D ("even") layers of the D-dimensional brickwork form hypercube clusters of size 2^D. The $(D+1)^{\mathrm{th}}$ layer, which is the first "odd" layer, then connects these hypercubes into rows of $2^D N^{1/D}$ sites. The remaining odd layers are then highly connected, so in the limit of large N their contribution can be ignored. This argument, given

in detail in Appendix F, shows that the $(\epsilon, t = 2)$-design depth of the D-dimensional brickwork is at most

$$d_* = 2D\frac{4N\log q + \log(1/\epsilon)}{\log[(1-\{1-\exp[-1/2C(q,2)]\}^2)^{-1}]}. \quad (68)$$

It seems likely that this argument could be improved to give a bound that actually decreases with D, since the dominant $(D+1)^{\text{th}}$-layer cluster-merging graph has a relatively simple structure. It may also be possible to extend such techniques to larger t and for other extensively connected architectures.

IX. RELATIONSHIP TO ARCHITECTURES THAT SCRAMBLE IN $O(\log N)$

This work establishes an $O(N)$ bound on the approximate t-design depth for any sequence of architectures with $\bar{\ell}$ independent of N. This is essentially a statement about worst-case architectures, since it is an upper bound over all possible architectures with any given $\bar{\ell}$. Other recent works have focused on engineering best-case architectures that converge especially quickly [26,27]. In this context, the fastest-known sequences of architectures have approximate t-design depths of at most $O(\log N)$.

The architectures constructed by Schuster *et al.* [26] have a few interesting properties. The main construction is of an architecture that is almost the 1D brickwork, but with certain gates removed so that $\bar{\ell} = O(\log N)$. Although their proof does not apply to the actual 1D brickwork, the extreme similarity of the two circuits suggests that the approximate t-design depth of the 1D brickwork architecture might also be $O(\log N)$. In addition, Schuster *et al.* [26] constructed architectures on arbitrary connectivity graphs with the same behavior. However, these architectures are qualitatively different from the generic architectures we study. They cannot be built with only Haar-random unitaries; they instead use SWAP gates to build an essentially 1D architecture on an arbitrary graph.

This raises an interesting question: do all reasonable sequences of architectures reach an approximate t-design at a rate bounded by $O(\log N)$, or is our $O(N)$ bound tight? In Sec. VIII A, we gave a heuristic argument that we should expect every connected architecture to scramble at least as quickly as the 1D brickwork architecture, in which case there would exist an $O(\log N)$ general architecture bound. This would be a significant advance if it could be proven. Another possibility is that there are multiple families of architectures, some of which scramble in depth $O(N)$ and some of which scramble in depth $O(\log N)$. In this case an understanding of the dividing line between these two qualitatively different classes of circuits would have important implications.

X. CONCLUSION

We show that bounds on approximate t-designs for 1D brickwork architectures imply bounds for general architectures. This process not only gives us an immediate bound linear in N for the tth moments of all sufficiently well-connected architectures, but allows us to convert any improved 1D bounds into bounds on generic structures. We also show that our bounds can be extended by an averaging procedure to an implicit bound for nondeterministic architectures.

Any architecture consisting of ℓ-layer connected blocks of $O(1)$ depth can be bounded this way. So our result implies that any sufficiently regularly connected circuit ensemble approximates global information scrambling in at most linear depth, albeit with exponential dependence on the connection frequency.

We conjecture that this bound can be tightened to one that depends on the circuit architecture only via the *connection count*. This suggests that rapid scrambling is inescapable for any sufficiently well-connected architecture.

ACKNOWLEDGMENTS

We acknowledge useful conversations with Nicholas Hunter-Jones. D.B. thanks Felix Leditzky for suggesting relevant literature. This material is based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers (B.K.C., J.A., D.B., B.F., S.G.) and work partially supported by the National Science Foundation under Grant No. CCF-1005 2044923 (CAREER). B.F. and S.G. acknowledge support 1006 from AFOSR (FA9550-21-1-0008) and DOE QuantISED Grant DE-SC0020360. We also acknowledge funding by the US Department of Energy Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing Program (C.K., F.T.C.); ARO W911NF-23-1-0077, NSF-OMA-2016136, and NSF Phy-1818914 (S.L.); and an NSF GRF (J.S.).

APPENDIX A: PROPERTIES OF PROJECTOR PRODUCTS

In this appendix we prove some general results about products of orthogonal projection operators. Consider some set of subspaces X_i, $i \in \{1,\ldots,n\}$, of a Hilbert space. Define P_i to be the orthogonal projector on to X_i and

$$T = P_n \cdots P_2 P_1.$$

In general, we are interested in understanding the singular value spectrum of T.

1. Structure of the unit eigenspace

Lemma 9. The left unit eigenspace, right unit eigenspace, and unit singular value space of T are all $\bigcap_{i=1}^{n} X_i$.

Proof. Let Y be the unit eigenspace of T. A projector is norm nonincreasing (i.e., $\|P_i\|_\infty = 1$). Furthermore, it acts as the identity on any vector whose norm it does not decrease. It follows that a unit eigenvector of T must be a unit eigenvector of each of the P_i. It is easy to see that the converse also holds, so $Y = \bigcap_{i=1}^n X_i$. This argument works the same from the left and the right, so the left and right unit eigenspaces are the same.

Now consider the singular value spaces of T. These are the square roots of the eigenvalues of $T^\dagger T$. Since the left and right eigenspaces are the same, they are in the unit eigenspace of $T^\dagger T$. Since both $T^\dagger$ and T are norm nonincreasing, they must be the whole unit eigenspace. So the unit singular value space of T is exactly the unit eigenspace. ■

Lemma 10. The unit eigenspace of T is orthogonal to all other eigenstates of T and all other eigenstates of $T^{\dagger k}T^k$, and remains orthogonal no matter how many factors of T or $T^\dagger$ are applied to the other eigenstate.

Proof. From Lemma 9 we know that the left and right unit eigenspaces of T are the same. Let u be a unit eigenvector and v a right eigenvector with eigenvalue $\lambda < 1$. We can compute

$$u^\dagger v = u^\dagger T v = \lambda u^\dagger v, \tag{A1}$$

which can hold only if $u^\dagger v = 0$. So Y is orthogonal to every nonunit eigenspace of T.

Since $T^{\dagger k}T^k$ is Hermitian, its eigenstates of different eigenvalues are automatically orthogonal, so $u^\dagger v = 0$ for any subunit eigenstates v of $T^{\dagger k}T^k$.

These proofs still hold if we apply extra factors of T or $T^\dagger$ to v, because we can freely absorb these extra factors into u. ■

2. Bound from layer-restricted subleading singular values

Theorem 13. Let Q_i to be the orthogonal projector on to $\bigcap_{j=1}^i X_j$. Let ϑ_i be the largest nonunit singular value of $P_i Q_{i-1}$. Let s_* be the largest nonunit singular value of T. Then we have the bound

$$s_*^2 \le 1 - \prod_{i=2}^n (1 - \vartheta_i^2). \tag{A2}$$

Proof. This is Theorem 7 of the main text. We prove Eq. (A2) by induction. Let $T_i = P_i \cdots P_1$ so that $T = T_n$. Suppose that T_{i-1} satisfies Eq. (A2). We prove that $T_i = P_i T_{i-1}$ also satisfies Eq. (A2).

Let Y_i be the unit eigenspace of T_i. Let v be a unit vector such that $\|T_i v\| = s_*(T_i)$. We may take an orthogonal decomposition $v = v_1 + v_2$, where $v_1 \in Y_{i-1}$ and $v_2 \perp Y_{i-1}$. We wish to compute

$$s_*(T_i) = \|P_i T_{i-1} v\| = \|P_i(v_1 + T_{i-1} v_2)\|.$$

Let θ, ϕ be the angles between $T_i v$ and $v_1, T_{i-1} v_2$, respectively. Note that $\|T_{i-1} v_2\| \le s(T_{i-1})\|v_2\|$. Then

$$\|P_i(v_1 + T_{i-1} v_2)\| \le \|v_1\| \cos\theta + s_*(T_{i-1})\|v_2\| \cos\phi. \tag{A3}$$

We next wish to optimize this bound over $\|v_1\|, \|v_2\|, \theta, \phi$ to obtain an unconditional bound. Our first constraint is $\|v_1\|^2 + \|v_2\|^2 = 1$, so after optimizing over the norms we find that

$$\|P_i(v_1 + T_{i-1} v_2)\| \le \sqrt{\cos^2\theta + s_*(T_{i-1})^2 \cos^2\phi} \equiv f(\theta, \phi). \tag{A4}$$

Now we optimize over the angles. By their definitions we must have $0 \le \theta \le \pi/2$ and likewise for ϕ. Note that f is a monotonically decreasing function of both θ and ϕ in this region, so the maximum will be attained somewhere on the boundary of the feasible set.

We also have some additional constraints. Since $v_1 \perp T_{i-1} v_2$, we must have $\theta + \phi \ge \pi/2$. And we know that the angle between v_1 and P_i must be at least $\cos^{-1}\vartheta_i$ by the definition of ϑ_i.

The Pareto frontier where both θ and ϕ are as small as possible lies along $\theta + \phi = \pi/2$, so the optimum must be somewhere on this line. On this line we have $\cos^2\phi = 1 - \cos^2\theta$, so

$$f(\theta) = \sqrt{[1 - s_*(T_{i-1})^2]\cos^2\theta + s_*(T_{i-1})^2}. \tag{A5}$$

Since $s_*(T_{i-1}) < 1$, this is again a decreasing function of θ, so it attains its maximum when θ is minimized, i.e., $\theta = \cos^{-1}\vartheta_i$. We thus find that

$$s_*(T_i) \le \sqrt{[1 - s_*(T_{i-1})^2]\vartheta_i^2 + s_*(T_{i-1})^2}. \tag{A6}$$

By assumption, $s_*(T_{i-1})^2 \le 1 - \prod_{j=2}^{i-1}(1 - \vartheta_j^2)$ and so

$$\begin{aligned} s_*(T_i)^2 &\le \vartheta_i^2 + (1 - \vartheta_i)^2\left(1 - \prod_{j=2}^{i-1}(1 - \vartheta_j^2)\right) \\ &\le 1 - \prod_{j=2}^{i}(1 - \vartheta_j^2), \end{aligned} \tag{A7}$$

completing the induction. Finally, for the base case, T_1 is a single projector, so $s_*(T_1) = 0$. ■

Previous works on approximate t-designs have often used the detectability lemma to relate the transfer matrix

to a frustration-free Hamiltonian [22,23]. The role of Theorem 7 in our proof is analogous to that of the detectability lemma in these works. The detectability lemma is based on counting the number of projectors that do not commute, whereas here the ϑ_i in some sense quantify the amount of noncommutativeness.

3. Other bounds on subleading singular values

Lemma 11. Let s_* be the largest nonunit singular value of T. Let λ_* be the nonunit eigenvalue of T with the largest magnitude. Then

$$|\lambda_*| \leq s_*. \tag{A8}$$

Proof. We may write

$$s_* = \max_{\{v|v\perp Z\}} \frac{\|Tv\|}{\|v\|}. \tag{A9}$$

Let v be an eigenvector of T with eigenvalue $\lambda \neq 1$. From before, $v \perp Z$, and, clearly, $\|Tv\|/\|v\| = \lambda$, so we have $s_* \geq \lambda$. The lemma follows immediately. ■

Theorem 14. Consider two projectors P_1, P_2. Let λ_* be the largest nonunit eigenvalue of P_2P_1, and let s_* be the largest nonunit singular value. Then

$$s_* = \sqrt{\lambda_*}. \tag{A10}$$

Proof. Any eigenvector v_* corresponding to λ_* must lie in the unit eigenspace of P_2. Otherwise, the output $P_2P_1v_*$ will not be parallel to v_*. Also, if we take the vector $w_* \equiv P_1v_*$, by Lemma 10, this vector is orthogonal to the unit eigenspace of P_2P_1, and

$$\begin{aligned}\frac{\|P_2P_1w_*\|^2}{\|w_*\|^2} &= \frac{v_*^TP_1P_2^2P_1v_*}{v_*^TP_1v_*} \\ &= \frac{\lambda_*^2}{v_*^T(P_2P_1v_*)} \\ &= \lambda_*.\end{aligned} \tag{A11}$$

So the subleading singular value of P_2P_1 is at least $\sqrt{\lambda_*}$.

Now, take w_* to be the subleading eigenvector of $(P_2P_1)^\dagger(P_2P_1) = P_1P_2P_1$. By definition, the corresponding eigenvalue is s_*^2. Take $v_* = P_2P_1w_*$ (still orthogonal to the unit eigenspace of P_2P_1 by Lemma 10). Then

$$\begin{aligned}P_2P_1v_* &= P_2P_1^2P_2P_1w_* \\ &= s_*^2P_2P_1w_* \\ &= s_*^2v_*.\end{aligned} \tag{A12}$$

We see that the subleading eigenvalue of P_2P_1 is at least s_*^2. Combining these two inequalities establishes Theorem 14. ■

Lemma 12. Let T_i, $i \in \{1,\ldots,k\}$, be a sequence of projector products that all share the same unit eigenspace. Let $s_*^{(i)}$ be the largest nonunit singular value of T_i. Then the largest nonunit singular value of $\prod_{i=1}^k T_i$ is at most $\prod_i s_*^{(i)}$.

Proof. Let $[M]_1$ denote the unit singular value space of M, and let $M_j = \prod_{i=j}^k T_i$. We can write the subleading singular value of M_k as

$$s(M_j) = \max_{v\perp[M_j]_1} \frac{\|M_jv\|}{\|v\|}. \tag{A13}$$

Since the T_i all share a unit eigenspace, the unit eigenspace of M_j is the same as that of each T_i. But, by Lemma 9, the unit eigenspaces and unit singular value spaces of both T_i and M_j are the same, so

$$[M_j]_1 = [T_1]. \tag{A14}$$

We thus have

$$s(M_j) = \max_{v\perp[T]_1} \frac{\|M_jv\|}{\|v\|}. \tag{A15}$$

Furthermore,

$$\|M_jv\| = \|M_{j+1}T_jv)\| \leq S_j\|M_{j+1}v\|, \tag{A16}$$

and so

$$s(M_j) \leq s_*^{(j)}s(M_{j+1}). \tag{A17}$$

Induction then gives the desired bound. ■

4. Bounds on Frobenius norms

Lemma 13. Let T be a projector product with largest nonunit singular value s_*. Let d be the dimension of the space on which T acts, let m_1 be the dimension of the unit eigenspace of T, and let m_0 be the dimension of the zero eigenspace. Then

$$\|T^k\|_F^2 \leq m_1 + (d - m_1 - m_0)s_*^2. \tag{A18}$$

Proof. Let σ_i be all the singular values of T. We have $\sigma_i = 1$ for $i \in \{1,\ldots,m_1\}$, $\sigma_j \leq s_*$ for $i \in \{m_1 + 1,\ldots,d - m_0\}$, and $\sigma_k = 0$ for $i\{\in d - m_0 + 1,\ldots,d\}$. We compute

$$\|T\|_F^2 = \sum_i \sigma_i^2 \leq \sum_{i=1}^m 1 + \sum_{i=m+1}^{d-m_0} s_*^2, \tag{A19}$$

from which the result follows immediately. We can further tighten the bound by replacing d with the number of nonzero singular values. ■

Theorem 15. Let T_i, $i \in \{1,\ldots,k\}$, be a sequence of projector products that all have the same unit eigenspace. Let d be the dimension of the space on which T_i acts, let m_1 be the dimension of the unit eigenspace of T_i, and let m_0 be the dimension of the zero eigenspace of any particular T_i. Let $s_*^{(i)}$ be the largest nonunit singular value of T_i. Then

$$\left\| \prod_i T_i \right\|_F^2 \leq m_1 + (d - m_1 - m_0) \prod_{i=1}^{k} (s_*^{(i)})^2. \tag{A20}$$

Proof. This follows directly from the results of Lemmas 12 and 13. ■

Theorem 16. Let $T_i, s_*^{(i)}, d, m_1, m_0$ be as in Theorem 15. Let $R_1,\ldots,R_n$ be a sequence of projector products such that the unit eigenspace of T_1 is contained within the unit eigenspace of each R_i. Let M be the product of all of the T_i and all of the R_i in any ordering. Then

$$\|M\|_F^2 \leq m_1 + (d - m_1 - m_0) \prod_{i=1}^{k} (s_*^{(i)})^2. \tag{A21}$$

Proof. The R_i also preserve the unit eigenspace of T_i and are also norm nonincreasing, so the proof of Lemma 12 still goes through. We can then again use Lemma 13 to obtain the final formula. ■

Theorem 17. Let T be a projector product with largest nonunit singular value s_*. Let d, m_1, m_0 be as before. Then

$$\|T^k\|_F^2 \leq m_1 + (d - m_1 - m_0) s_*^{2k}. \tag{A22}$$

Proof. This is an immediate consequence of Theorem 15 in the case where the T_i are all the same. ■

APPENDIX B: PROOF OF THE 1D BRICKWORK SPECTRAL GAP FOR $t = 2$

For a string of N sites that we act a 1D brickwork architecture on, we can consider the nonorthogonal basis

$$|\vec{X}\rangle = \bigotimes_i |X_i\rangle_i, \tag{B1}$$

where X_i refers to one of the two $k = 2$ permutation states $|I\rangle$ (identity) or $|S\rangle$ (swap). This basis is complete on the image of any layer of the brickwork. We start with the following result.

Lemma 14. Take P_1 to be the projector into the unit eigenspace of T. Then there exists a depth-independent constant $c_{X'X}$ such that

$$\langle \vec{X}'|T^k|\vec{X}\rangle \leq \langle \vec{X}'|P_1|\vec{X}\rangle + c_{X'X}\lambda_1^k \tag{B2}$$

for every $k > \log(N)/[-\log(\lambda_1)] + 1$, where $\lambda_1 = [2q/(q^2+1)]^4$.

Proof. We use a domain-wall trajectory approach [21]. Each gate sends $|II\rangle \to |II\rangle$, $|SS\rangle \to |SS\rangle$, and the nonuniform $|IS\rangle, |SI\rangle \to q/(q^2+1)(|II\rangle + |SS\rangle)$. The point is that the transfer matrix sends each configuration into a sum of other configurations, depending on the positions of the I, S domain walls in the system. If a gate in the transfer matrix crosses a domain wall (and all domain walls will be crossed by a gate each layer after the first layer), it either moves left with weight $q/(q^2+1)$ or right with weight $q/(q^2+1)$. So we can represent the transfer matrix of k layers as a series of domain-wall trajectories with their accompanying weights. A domain-wall trajectory is a sequence of $\vec{X}^j$'s such that $\vec{X}^0 = \vec{X}$ and $\vec{X}^d = \vec{X}'$. Specifically, we are looking out for the domain walls in each layer $\vec{X}^j$, because the total number of domain walls never increases—each domain wall either moves around or annihilates with a neighbor, possibly eventually reaching the steady states $|I\rangle^N, |S\rangle^N$ with no domain walls. We have, for domain-wall trajectories γ with final state F_γ and weight $w(\gamma)$,

$$\langle \vec{X}'|T^k|\vec{X}\rangle = \sum_\gamma w(\gamma)\langle \vec{X}'|F_\gamma\rangle. \tag{B3}$$

We can categorize each domain-wall trajectory according to which domain walls annihilate before the end of the circuit and which remain to the end. We can separate the total domain-wall trajectory into the annihilating ("closed") domain walls γ_c from the "open" domain-wall trajectories γ_o that stay to the end. The weight of the whole domain-wall trajectory is just a product of the two components:

$$w(\gamma) = w(\gamma_o)w(\gamma_c). \tag{B4}$$

This completes the proof. ■

Moreover, we have the following result.

Lemma 15. The final state F_γ is determined solely by the open component of the trajectory.

Proof. If we look at one of the surviving domain walls, the I/S values to the immediate left or right of the domain wall must remain the same through the whole trajectory (otherwise, the domain wall would be annihilated). Then the rest of the final configuration can be determined by just flipping the sign when crossing every other domain wall. Because the signs to the immediate left and right of the domain wall were fixed by the original configuration $\vec{X}$, there is no global symmetry ambiguity either. ■

With this we can separate sum (B3) into a sum over the partitions p that specify which domain walls are closed and which remain open. The partition in turn specifies the set

of possible open and closed domain-wall trajectories, $\mathcal{O}(p)$ and $\mathcal{C}(p)$:

$$\begin{aligned}\langle\vec{X}'|T^k|\vec{X}\rangle &= \sum_p \sum_{\gamma_o\in\mathcal{O}(p)} \sum_{\gamma_c\in\mathcal{C}(p)} w(\gamma_o)w(\gamma_c)\langle\vec{X}'|F_\gamma\rangle \\ &= \sum_p \sum_{\gamma_o\in\mathcal{O}(p)} w(\gamma_o)\langle\vec{X}'|F_\gamma\rangle\Big(\sum_{\gamma_c\in\mathcal{C}(p)} w(\gamma_c)\Big).\end{aligned} \tag{B5}$$

Now we start to bound things.

Lemma 16. The sum $\sum_{\gamma_c\in\mathcal{C}(p)} w(\gamma_c)$ is bounded by the infinite depth limit, i.e., the sum of all the possible trajectories in an infinite depth circuit that start at the domain walls specified by p, but annihilate to either the uniform $|I\rangle$ or $|S\rangle$ states.

This is because $\mathcal{C}(p)$ is just a subset of all the possible closed trajectories, and increasing the depth on a trajectory that is already closed does not change its weight. The sum of weights in the infinite depth limit, in turn, has to be

$$\lim_{k\to\infty} |T^k|\vec{X}_C(p)\rangle| = |P_1|\vec{X}_C(p)\rangle|, \tag{B6}$$

where $|\vec{X}_C(p)\rangle$ is the initial configuration given by only the closed domain walls, and P_1 is the unit eigenspace projector. That is, it is the component of $|\vec{X}_C(p)\rangle$ in the unit eigenspace of T, because the unit eigenspace is all that survives in the infinite depth limit. Let us call this infinite depth limit $W(p)$. We also have the following result.

Lemma 17. It holds that

$$\sum_{\gamma_o\in\mathcal{O}(p)} w(\gamma_o) \le [2q/(q^2+1)]^{2N_o(k-1)}, \tag{B7}$$

where N_o is the number of open domain walls.

Proof. We build up $\sum_{\gamma_o\in\mathcal{O}(p)} w(\gamma_o)$ layer by layer. At each layer beyond the first, each open trajectory γ_o needs to move each of its N_o domain walls. Each domain wall has at most two possible directions to move in (and could be less than two if other domain walls are blocking the way). No matter the direction, the domain wall acquires a weight $q/(q^2+1)$ by moving. So the total weight of all the possible trajectories created by adding a layer onto γ_o is at most $w(\gamma_o)[2q/(q^2+1)]^{N_o}$. A transfer matrix is composed of two layers, so this is an extra factor of $[2q/(q^2+1)]^{2N_o}$ per transfer matrix applied. Extrapolating back to the first layer gives us the equation above. ■

With these two inequalities, we have

$$\begin{aligned}\langle\vec{X}'|T^k|\vec{X}\rangle &\le \sum_p W(p) \sum_{\gamma_o\in\mathcal{O}(p)} w(\gamma_o)\langle\vec{X}'|F_{\gamma_o}\rangle \\ &\le \sum_p W(p)x_F \sum_{\gamma_o\in\mathcal{O}(p)} w(\gamma_o) \\ &\le \sum_p W(p)x_F\lambda_1^{N_o(k-1)/2}\end{aligned} \tag{B8}$$

for $x_F \equiv \max_{\gamma_o\in\mathcal{O}(p)}\langle\vec{X}'|F_{\gamma_o}\rangle$ and $\lambda_1 = [2q/(q^2+1)]^4$. The right-hand side is now a weighted sum of terms that are exponential in depth, with base $\lambda_1^{N_o/2}$ dependent on the number of open trajectories in the partition. The gentlest exponential is the single partition where $N_o = 0$ and all the domain walls are closed—in that case $W(p)$ is the component of $|\vec{X}\rangle$ in the unit eigenspace and $[\max_{\gamma_o\in\mathcal{O}(p)}\langle\vec{X}'|F_{\gamma_o}\rangle]$ is the component of $|\vec{X}'\rangle$ in the unit eigenspace. This term is therefore bounded by $\lim_{k\to\infty}\langle\vec{X}'|T^k|\vec{X}\rangle$, which, by Lemma 2, is the same as $\langle\vec{X}'|P_1|\vec{X}\rangle$

In periodic boundary conditions, the rest of the partitions all have $N_o \ge 2$. Hence, they decay at a rate λ_1^k or faster. Specifically, we have

$$\langle\vec{X}'|T^k|\vec{X}\rangle \le \langle\vec{X}'|P_1|\vec{X}\rangle + c_{X'X}\lambda_1^k \tag{B9}$$

for some depth-independent constant $c_{X'X}$ (note that while $\max_{\gamma_o\in\mathcal{O}(p)}\langle\vec{X}'|F_{\gamma_o}\rangle$ is depth dependent, it is bounded above by 1). Now we can prove that $\lambda_1 = [2q/(q^2+1)]^4$ genuinely is the subleading eigenstate of T, using the following result.

Lemma 18. For a complete basis $|\vec{X}\rangle$, if $\langle\vec{X}'|T^k|\vec{X}\rangle \le \langle\vec{X}'|P_1|\vec{X}\rangle + c_{X'X}\lambda_1^k$ and $\langle\vec{X}'|T^k|\vec{X}\rangle \ge \langle\vec{X}'|P_1|\vec{X}\rangle$, then T^k has no eigenstate $|\psi_2\rangle$ with eigenvalue $\lambda_1 < \lambda_2 < 1$.

Proof. Suppose that there exists such an eigenstate $|\psi_2\rangle$. Then, because $|\vec{X}\rangle$ is complete, there must exist some $\vec{X}$ that has a nonzero component of $|\psi_2\rangle$ in its eigenstate decomposition of T:

$$|\vec{X}\rangle = a_1|\psi_1\rangle + a_2|\psi_2\rangle + \cdots, \qquad a_2 \ne 0. \tag{B10}$$

Moreover, there must exist some $\vec{X}'$ that has a nonzero overlap with $|\psi_2\rangle$, i.e., $\langle\vec{X}'|\psi_2\rangle \ne 0$. Then we have

$$\begin{aligned}\langle\vec{X}'|T^k|\vec{X}\rangle &= a_1 1^k\langle\vec{X}'|\psi_1\rangle + a_2\lambda_2^k\langle\vec{X}'|\psi_2\rangle + \cdots \\ &= 1^k\langle\vec{X}'|P_1|\vec{X}\rangle + a_2\lambda_2^k\langle\vec{X}'|\psi_2\rangle + \cdots.\end{aligned} \tag{B11}$$

In particular, because $\lambda_2 > \lambda_1$, for any constant $c_{X'X}$, there must be some k for which $\langle\vec{X}'|T^k|\vec{X}\rangle > \langle\vec{X}'|P_1|\vec{X}\rangle + c_1\lambda_1^k$. This is a contradiction of our assumption, so no such λ_2 can exist.

One caveat is that the eigenstate overlap $a_2\langle\vec{X}'|\psi_2\rangle$ could be negative instead. This is where the lower bound comes in—provided λ_2 is the largest subleading eigenvalue, there must also exist some k for which $a_2\lambda_2^k\langle\vec{X}'|\psi_2\rangle$ overtakes every nonunit term in the sum and the lower bound is violated instead. This lower bound is naturally satisfied in our case because P_1 is the infinite depth limit of T^k, but T is contractive, so adding more layers to $\langle\vec{X}'|T^k|\vec{X}\rangle$ always decreases its norm. Because it is a sum of positive trajectories, it is also positive, so the value of $\langle\vec{X}'|T^k|\vec{X}\rangle$ cannot increase as it goes to P_1. ■

Once we have $\lambda_1 = [2q/(q^2+1)]^4$, by Lemma 11, $s_* = [2q/(q^2+1)]^2$. Therefore, our trace decays at a rate $C(q,2) = 1/2\log(1/s_*) = 1/4\log[(q^2+1)/2q]$.

APPENDIX C: MAPPING INCOMPLETE TO COMPLETE LAYERS FOR NONINTEGER $\sqrt{q}$

If $\sqrt{q}$ is not an integer, we cannot draw Haar random $\sqrt{q}\times\sqrt{q}$ unitaries. There is thus no such thing as the 1D brickwork circuit ensemble. So how can we apply the site-splitting trick used in the proof of Theorem 2? Instead of defining $s_{1D}(q)$ to be the subleading singular value of the transfer matrix corresponding to some underlying circuit ensemble, we define the transfer matrix directly.

The first step is to rephrase our site-splitting strategy from something done in the quantum circuits to something done at the level of transfer matrices. Consider a set of vector spaces labeled by a positive real valued $r > 1$, i.e.,

$$X_r = \mathrm{span}\{|\sigma\rangle_r \,|\, \sigma \in S_t\}, \tag{C1}$$

where each vector space is equipped with a basis $|\sigma\rangle_{X_r}$ and an inner product

$$\langle\sigma|\tau\rangle_r = r^{|\sigma\tau^{-1}|}. \tag{C2}$$

Here $|\sigma|$ is the length of the cycle structure of σ. Note that this inner product is positive semidefinite only for integer r.

Now, consider the mapping $V\colon X_r \to X_{\sqrt{r}}^{\otimes 2}$ defined on basis elements by

$$V|\sigma\rangle_r = |\sigma\rangle_{\sqrt{r}}^{\otimes 2}. \tag{C3}$$

We first show that V is an isometry. Compute

$$\langle V|\sigma\rangle, V|\tau\rangle\rangle_{\sqrt{r}} = \langle\sigma\sigma|\tau\tau\rangle_{\sqrt{r}} = (\sqrt{r}^{|\sigma\tau^{-1}|})^2 \tag{C4}$$

and

$$\langle\sigma|\tau\rangle_r = r^{|\sigma\tau^{-1}|}, \tag{C5}$$

which are the same. This implies that the restriction of the metric on $X_{\sqrt{r}}^{\otimes 2}$ to the image of V is positive semidefinite (for integer r).

For each X_r, let us define the k-site gate $G_k^{(X_r)}\colon X_r^{\otimes k} \to X_r^{\otimes k}$ as the following projector onto the span of $\{|\sigma\rangle_r^{\otimes k} \,|\, \sigma \in S_t\}$:

$$G_k^{(X_r)} = \sum_{\sigma,\tau} |\tau\rangle^{\otimes k} \mathrm{Wg}(r^k)_{\sigma\tau} \langle\sigma|^{\otimes k} \tag{C6}$$

with $\mathrm{Wg}(q)$ the Moore-Penrose pseudoinverse of the metric $g(q)_{\sigma\tau} = \langle\sigma|\tau\rangle_r$. This formula reproduces the usual k-site gate when r is an integer. We now observe the following.

Lemma 19. If we replace every $G_k^{X_r}$ in the transfer matrix with $G_{2k}^{X_{\sqrt{r}}}$, the singular values do not change.

Proof. Let us define a map $V^\dagger$ by

$$V^\dagger = \sum_{\sigma\tau} |\sigma\rangle_r \,\mathrm{Wg}(r)_{\sigma\tau}\, \langle\tau\tau|_{\sqrt{r}}, \tag{C7}$$

so that

$$V^\dagger\, |\tau\tau\rangle_{\sqrt{r}} = |\tau\rangle_r .$$

This map is an adjoint of V on the image of V, i.e.,

$$\langle\sigma|_r\, V^\dagger\, |\tau\tau\rangle_{\sqrt{r}} = \langle\tau\tau|_{\sqrt{r}}\, V|\sigma\rangle_r \tag{C8}$$

for all τ, σ.

From the definition of V, we see that

$$G_k^{X_r} = V^{\dagger\otimes k} G_{2k}^{X_{\sqrt{r}}} V^{\otimes k}. \tag{C9}$$

We may thus rewrite the transfer matrix T by replacing $G_k^{X_r}$ with the above expression for each gate. Furthermore, $VV^\dagger = G_1^{X_{\sqrt{r}}}$. Factors of G_1 may be absorbed into G_k from either side, which means in particular that

$$G_{2k}^{X_{\sqrt{r}}} = (VV^\dagger)^{\otimes k} G_{2k}^{X_{\sqrt{r}}} = G_{2k}^{X_{\sqrt{r}}} (VV^\dagger)^{\otimes k}. \tag{C10}$$

We may thus pair up the copies of V that appear on internal legs of the transfer matrix and absorb them into the adjacent $G_{2k}^{X_{\sqrt{r}}}$. The singular values of T are the nonzero eigenvalues of $T^\dagger T$, so the copies of V that appear on input legs can be cycled to the output and canceled against the corresponding copies of $V^\dagger$ without changing the singular values. ■

If we return to our original transfer matrix, we see that we can identify each site as a member of X_q, and each two-site gate as a copy of $G_2^{X_q}$. Moreover, we can freely apply copies of $G_1^{X_q}$ (the averaged one-site gate) to any site wherever we want, as it is just the identity on that vector space. Now we apply isomorphism V to map each site from an

element of X_q to an element of the doubled space $X_{\sqrt{q}}^{\otimes 2}$, converting each site into a pair of sites, or *twinned sites*, in the process. From Lemma 19, we see that the transfer matrix can be rewritten in this new vector space, without changing any singular values, by replacing $G_k^{X_q}$ with $G_{2k}^{X_{\sqrt{q}}}$ everywhere it appears.

In the particular case of an incomplete circuit, we obtain a transfer matrix consisting of $G_1^{X_q}$ and $G_2^{X_q}$. We then apply the isomorphism described above to map $G_1^{X_q} \to G_2^{X_{\sqrt{q}}}$ and $G_2^{X_q} \to G_4^{X_{\sqrt{q}}}$.

The former map sends one-site gates to two-site gates; these correspond to edges across two twinned sites. The first layer of the original circuit involved a one- or two-site gate acting on every site, so the split pair sites are always joined back into the same cluster by the first layer. The new four-site gates are harder to express in the cluster-merging picture. However, we know that a mini circuit of two-site gates $G_2^{X_{\sqrt{q}}}$, applied in a way that connects all the sites consistently, will approach $G_4^{X_{\sqrt{q}}}$ in the limit of infinite layers. In particular, we can replace each four-site gate $G_4^{X_{\sqrt{q}}}[a_1, a_2, b_1, b_2]$ acting over sites $a_1, \ldots, b_2$ (where a_1 and a_2 are twinned, and so are b_1 and b_2) with an arbitrarily large "Jenga tower" of gates

$$(G_2^{X_{\sqrt{q}}}[a_1, a_2] \otimes G_2^{X_{\sqrt{q}}}[b_1, b_2])(G_2^{X_{\sqrt{q}}}[a_1, b_1] \otimes G_2^{X_{\sqrt{q}}}[a_2, b_2])$$

repeated over and over again. In the cluster-merging picture, only the bottom layer of this tower is capable of joining distinct clusters together. The other layers either join twinned sites that were part of the same cluster to begin with, or are copies of the bottom layer. So every layer above the first has a completely disconnected cluster-merging graph, which corresponds to $\mathfrak{z} = 0$. These layers do not contribute to the right-hand side of Theorem 7. So each four-site gate can be replaced with two two-site gates connecting nontwinned members together without increasing the subleading singular value of the transfer matrix. This process is illustrated in Fig. 6.

We see that we can replace the cluster-merging picture with one where every site is replaced by twins on $X_{\sqrt{q}}$, and every two-site gate is replaced by a pair of two-site gates between nontwinned sites. Moreover, we can freely apply two-site gates between any twinned members where no gate was being originally applied. So all sites in the new picture have two-site gates acting on them; the layer is complete. Since all twinned sites belong to the same cluster, all clusters are of even size. These were the two conditions required to draw a Eulerian cycle on each connected component of the graph. We can now use Lemma 5 to reduce this graph to a periodic 1D brickwork architecture composed of gates $G_2^{X_{\sqrt{q}}}$ acting on the space $X_{\sqrt{q}}^{\otimes 2N}$. Any previous work that has found a bound for the $\sqrt{q}$ brickwork therefore imposes a bound on the layer-restricted subleading singular values of arbitrary cluster-merging graphs.

APPENDIX D: GRAPH-SPLITTING BOUNDS FOR INCOMPLETE LAYERS

In this appendix we consider strategies for bounding cluster-merging graphs that originate from incomplete layers, without resorting to 1D brickwork bounds on noninteger q.

1. Analytical bound in terms of the node degree

By removing edges according to Lemma 8, we can reduce any connected cluster-merging graph into a spanning tree. Let d be the maximum degree of the tree. To bound this tree's singular value, we use the following tool for incomplete graphs.

Lemma 20. The edges of a cluster-merging graph can be split into separate layers without lowering the subleading singular values.

Proof. This follows from a simple reinterpretation of which gates belong to the same layer. We know that the gates corresponding to each edge in the cluster-merging graph commute with each other. Therefore, we are allowed to choose which gates to apply first. Splitting a layer merely means choosing a subset of gates to apply first, then choosing another subset to apply in the next new layer, and so on. This does not lower the subleading singular values because it does not change the gates at all. ■

Note that splitting a layer will increase the block width ℓ, so our bound on the subleading singular value of the overall transfer matrix via Theorem 13 will get looser when a layer is split. Nonetheless, this will be useful for separating the degree-d tree graph into reducible parts.

Theorem 18. A cluster-merging graph of size N with a spanning tree of maximum degree d can be decomposed into at most

$$2\min(\lceil d/2 \rceil, 6)\lceil \log_2(N) \rceil$$

layers, such that each layer is reducible to the 1D brickwork.

We prove this statement starting with the following similar result.

Lemma 21. A cluster-merging graph of size N with a spanning tree of maximum degree d can be decomposed into at most $\min(\lceil d/2 \rceil, 6)\lceil \log_2(N) \rceil$ layers, such that each layer is made up of isolated strings.

Call $L(N)$ the largest possible number of layers a tree of size N needs to be decomposed into to satisfy this. We prove by induction that $L(N) \leq \min(\lceil d/2 \rceil, 6)\lceil \log_2(N) \rceil$. We have $L(2) = 1$ because two connected sites automatically form an isolated string. It remains to show the following result.

Lemma 22. It holds that $L(N) \leq L(\lfloor N/2 \rfloor) + \min(\lceil d/2 \rceil, 6)$.

Proof. Take an arbitrary root node C_1 in our tree and construct a path $C_1 C_2 \cdots C_r$ starting from that root. At each step $i-1$ in the path, we choose the child C_i of parent C_{i-1} with the most nodes in its subtree. Call the other child nodes $b_i^{(1)}, b_i^{(2)}, \ldots, b_i^{(g)}$, where $g \leq d-2$.

For a given path node C_{i-1}, each nonpath child $b_i^{(1)}, b_i^{(2)}, \ldots, b_i^{(g)}$ has at most $\lfloor N/2 \rfloor$ nodes in its subtree. We can therefore recursively decompose each subtree in parallel in at most $L(\lfloor N/2 \rfloor)$ layers. After these layers, each subtree of $b_i^{(j)}$ has been combined into $b_i^{(j)}$ to form one cluster, $B_i^{(j)}$.

We then need to combine all g child nodes $B_i^{(j)}$ into C_{i-1}. To accomplish this in an efficient way, we use a method for efficient contraction of high-degree nodes. ■

Lemma 23. We can combine a node with g of its external neighbors using at most four layers of isolated strings.

Proof. To accomplish this, we combine our layer splitting process with a cluster splitting process. The original node has a degree of g, so it must contain at least g sites. We first split the layer, with $\lfloor g/2 \rfloor$ of the external neighbors on the bottom layer. This gives us a node with g sites and $\lfloor g/2 \rfloor$ edges. We then use Lemma 7 to split this node into a string of nodes such that each node connects to two of the external neighbors (as shown by the first step in Fig. 7). Each node therefore has two external neighbors and at most two new internal edges connecting it to the other nodes in the string, so can be made with four sites. The exceptions are the end nodes that need one less site, and if $\lfloor g/2 \rfloor$ is odd, one node will have one less neighbor and site. In total, at most $4\lceil \frac{1}{2}\lfloor g/2 \rfloor \rceil - 2 \leq g$ sites are required, so there are enough unoccupied sites in the original node to add the necessary new edges for this splitting.

Once we have split the original node in this way, we spend two layers to combine all the neighbors and nodes together. In the first layer, each node in the string combines with its two external neighbors, as together they form a string of length 3 (Fig. 7, second step). With all the neighbors absorbed, the string can be recombined into the original node in the second layer (Fig. 7, third step).

We then spend two more layers to contract the remaining $g - \lfloor g/2 \rfloor$ neighbors. This is either $\lfloor g/2 \rfloor$ or $\lfloor g/2 \rfloor + 1$ depending on the parity of g. Even if it is the latter, we will have enough sites to contract all the neighbors, due to the $\geq \lfloor g/2 \rfloor$ extra sites the root node gained by absorbing its previous neighbors. Therefore, all neighbors can be contracted in four layers. ■

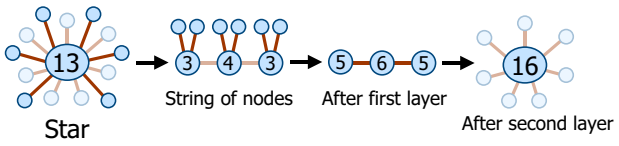

FIG. 7. The star contraction of Lemma 23. We start with a star with g external edges, meaning that its site count (the numerical label of the node) is at least g. In the first pair of layers, half of the edges are considered, then the star node is split into a string of nodes, each with two external edges. In the first layer, each node contracts with its two external edges in a three-site string (note that the site count on each node is just the minimum number of sites required to have the correct degree on the node; extra sites can be placed in whatever node we want). In the second layer, the resulting string is contracted back into a single node, leaving a star with half the edges remaining. The process is then repeated for the other half of the edges.

Note that Lemma 23 is inefficient if the number of leaves is 6 or less. This is because we can also contract g leaves in $\lceil g/2 \rceil$ layers, by selecting one or two leaves in each layer and combining with C_{i-1} into a string of length 2 or 3. So the number of layers we have to spend is $\min(\lceil g/2 \rceil, 4) \leq \min(\lceil d/2 \rceil - 1, 4)$ overall.

By performing all of these operations in parallel for each path node, we have spent at most $L(\lfloor N/2 \rfloor) + \min(\lceil d/2 \rceil - 1, 5)$ layers to combine each path node's nonpath children into the path node. After these layers, the cluster-merging picture of the next layer only consists of the path nodes, which can be combined together with a single, extra, layer. This completes the proof of Lemma 22.

We have seen that we can decompose the size-N cluster into $L(N)$ layers such that the cluster merging graph of each layer consists of isolated strings. Each string can now be split into an open 1D brickwork architecture, as long as each internal node is of even size (endpoint nodes can be either odd or even), using Lemma 5. We call such strings *brickwork compatible*. For strings containing nodes that are not even and are not on the ends, we can use the following result.

Lemma 24. Any layer of isolated strings can be split into two layers of brickwork-compatible strings.

Proof. We wish to split a string into two layers such that both layers have cluster-merging graphs consisting entirely of brickwork-compatible strings. Let k be the number of odd-sized nodes in the string and number only the odd nodes $1, \ldots, k$.

First suppose that k is even. Remove the left-hand edge, if any, of odd nodes $1, 3, \ldots, k-1$ and the right-hand edge of odd nodes $2, 4, \ldots, k$. Now we have split the string into at least $k/2$ and at most $k/2+2$ substrings, each of which either has odd nodes only at both endpoints, or has no odd nodes at all. Since the odd nodes were paired up, contracting the first layer gives only even nodes for the second layer. So both layers are brickwork compatible.

Now suppose that k is odd. Remove the left-hand edge, if any, of odd nodes $1, 3, \ldots, k-1$ and the right-hand edge of odd nodes $2, 4, \ldots, k-1$. Every substring has two odd endpoints except the first, which may have zero, and the last, which has exactly one. After contracting, the second layer has all even nodes except for one. But this one odd node is the right endpoint of the second-layer string, so the second layer is still brickwork compatible. ■

Lemma 24 shows that we can make a set of layers of isolated strings brickwork compatible with a splitting scheme that at most doubles the number of layers. Combining this with Lemma 21 completes the proof of Theorem 18.

2. Log log bound on arbitrary graphs

Theorem 19. A cluster-merging graph of size N can be decomposed into at most

$$8\lceil \log_2 \lfloor \log_2 (N+1) \rfloor \rceil + 2$$

layers, such that each layer is reducible to the 1D brickwork.

Given Lemma 24, it is sufficient to prove that we can decompose the graph into

$$4\lceil \log_2 \lfloor \log_2 (N+1) \rfloor \rceil + 1$$

layers of isolated strings.

We begin with a tree of at most N sites [Fig. 8(a), left]. The first layer of edges we apply will contract substrings of the tree such that the resulting second-layer tree [Fig. 8(a), right] has depth $O(\log N)$. Choose any root node, which also specifies a direction in the tree from its root down to the leaves. Then label each node by its *subtree weight*—the number of nodes in the subtree starting from that node. The root node has subtree weight N, leaves have subtree weight 1, and so on. We add to the first layer the *maximally weighted path* starting from the root node, i.e., a path from root to leaf that always selects the child node with the highest subtree weight. For every node that neighbors this path (excluding nodes that are part of the path), we take its subtree, add the maximally weighted path of that subtree to the first layer, and repeat the process recursively. After the end of this process, each node is part of exactly one maximally weighted path, so the first layer consists of disconnected strings. We can also label each maximally

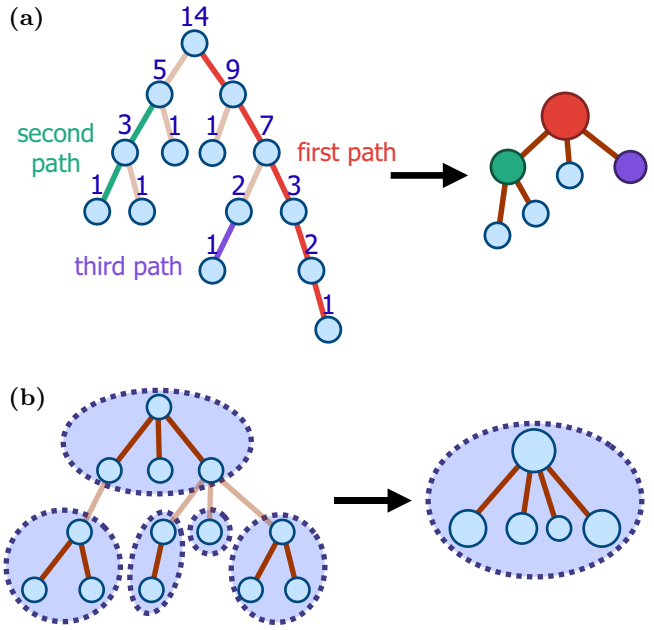

FIG. 8. Algorithm for reducing a general tree in $O(\log \log N)$ layers. (a) The first layer of the algorithm, where we contract the maximally weighted paths. Each path chooses the child node with the highest subtree weights (blue labels above each node), and starts from nonpath children of nodes in higher paths. (b) Contraction of a balanced tree into a point in $O(\log \log N)$ steps. At each step, the nodes at depths $2k$ and $2k+1$ are merged together through the star contraction algorithm (each star is indicated by a blue dotted circle).

weighted path by a subtree weight, which in this case is the subtree weight of the node at the start of the path.

After contracting all the paths added to the first layer this way, we are left with a tree where each node corresponds to one of the maximally weighted paths. We make the following claim.

Lemma 25. The height of the new tree is at most $\lfloor \log_2 (N+1) \rfloor \rceil - 1$.

Proof. Suppose that we start from a node of the new tree and jump down to one of its children. This corresponds to jumping from one maximally weighted path on the original tree to one of the paths below it.

Going back to the original tree, consider a maximally weighted path of a particular single-root subtree consisting of n nodes. The subtree weight of the root is, by definition, n, and so is the subtree weight of the path itself. Then any node that is one edge away from the maximally weighted path (i.e., it is not part of the maximally weighted path, but its parent is) has a subtree weight bounded above by $\lfloor (n-1)/2 \rfloor$. We can prove this by contradiction—it cannot be the root node, and the sum of the subtree weights of it and all its sibling nodes must be $\leq n-1$. So, if it had a subtree weight greater than $(n-1)/2$, it would have had a

larger subtree than all of its siblings, and therefore would be part of the maximally weighted path after all.

Every time we move from a maximally weighted path to a lower one, our subtree weight roughly halves; specifically, it goes from n to at most $\lfloor (n-1)/2 \rfloor$. If we repeat this process, we run out of nodes after $\lfloor \log_2(n+1) \rfloor - 1$ steps. Hence, the height of the whole second-layer tree must be bounded by $\lfloor \log_2(N+1) \rfloor \rceil - 1$. ■

Now we reduce the second-layer tree with the following lemma.

Lemma 26. A tree of height h can be decomposed into $4\lceil \log_2(h+1) \rceil$ layers of isolated strings.

Proof. We decomposed the tree recursively. At each step, we take all the nodes that are at an even depth (including the root node at depth 0). We can then use Lemma 23 on each even-depth node to absorb all their children in parallel, in at most four layers [Fig. 8(b), left]. This way, after four layers, we have removed all the nodes that were at an odd depth. Hence, we are left with a tree of height $\lceil (h+1)/2 \rceil - 1$ [Fig. 8(b), right]. We repeat this process until we are left with a tree of height 0, i.e., just the root node. This takes $\lceil \log_2(h+1) \rceil$ steps, or $4\lceil \log_2(h+1) \rceil$ layers, in total. ■

Hence, we can spend one layer to reduce an arbitrary tree to a tree of height $\lfloor \log_2(N+1) \rfloor - 1$, then spend $4\lceil \lfloor \log_2(N+1) \rfloor \rceil$ layers to reduce the new tree to a single node, which proves Theorem 19.

APPENDIX E: NUMERICAL EVIDENCE FOR CONJECTURES

1. Evidence for t independence of s_{1D}

Figures 9 and 10 present numerical evidence for Eqs. (19) and (64), respectively. We calculate $s_{1D}(N,q,t)$ with

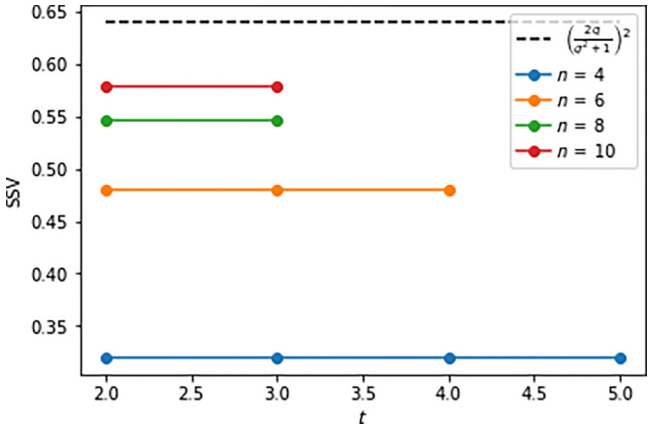

FIG. 9. Subleading singular values of the 1D brickwork architecture with periodic boundary conditions and $q = 2$, calculated via `svds`.

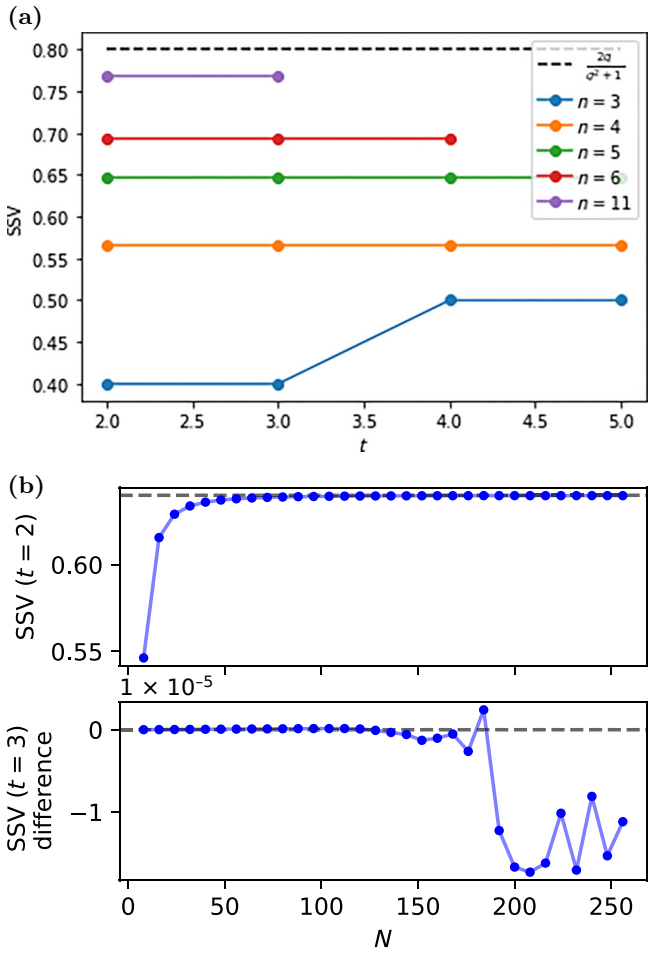

FIG. 10. (a) Subleading singular values of the 1D brickwork architecture with open boundary conditions and $q = 2$, calculated via `svds`. (b) DMRG results for the open 1D brickwork subleading singular value $s_{1D,open}(N,q,t=2)$ versus N (*upper*) and its difference from the $t = 3$ SSV (*lower*). Note the vertical axis scale of 10^{-5} in the lower panel.

both open and closed boundary conditions for $q = 2$ and several values of N and t.

Calculations are done using two different approaches. The SciPy library's `svds` function, which is based on the Lanczos algorithm, allows us to find eigenvectors with arbitrary precision for small N. Here the absolute error tolerance is set to 10^{-4}.

Because of the transfer matrix's natural representation as a tensor network, and hence a matrix product operator, we can also use the density matrix renormalization group (DMRG) algorithm to approximate eigenvectors for much larger N.

The accuracy of the DMRG algorithm depends on the entanglement structure of the true eigenvector, quantified by its bond dimension. Here we choose a singular value cutoff of 10^{-12} and a maximum bond dimension of 800, which is well above the $t = 2$ or $t = 3$ required bond dimension of approximately 12.

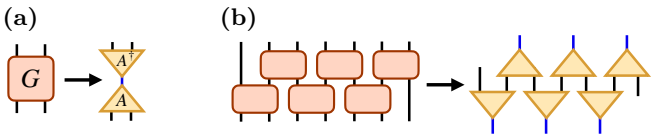

FIG. 11. (a) Decomposition of a single gate into two halves using SVD. The inner index (blue) can be treated as effective physical indices with dimension $t!$. (b) Decomposition of two layers of a 1D brickwork architecture using this splitting process. The bottom-most and top-most halves of the original brickwork architecture are set aside for now, making this equivalent to one brickwork layer in thickness. Through this process, the physical space has reduced from N legs of size q^{2t} to $\lceil N/2 \rceil$ legs of size $t!$.

Both methods are computationally intractable in the original representation of G, which has four legs, each of dimension q^{2t}. However, in this representation G is very sparse. For the brickwork circuit, we can use singular value decompositions to compress G to a tensor with three legs, each of dimension $t!$. The resulting tensor is an orthogonal projector from S_t to S_t^2 under the metric induced by the Weingarten functions. This compression leaves the nonzero singular values of the whole circuit unchanged. The compressed circuit is illustrated in Fig. 11.

2. Evidence for Conjecture 1

To search for violations of Conjecture 1, we used a simulated annealing algorithm that attempts to maximize the subleading singular value over the set of architectures on a fixed number of sites. Figure 12 shows the results.

Each proposed move was a set of n edge additions or deletions, with n drawn from a geometric distribution with mean 1. Additions and deletions were equally likely, with additions drawn uniformly from the all-to-all graph and deletions drawn uniformly from the set of existing edges.

The objective function was the third-largest singular value, with starting temperature set automatically based on the distribution of singular values over a small sample of connected architectures with edges drawn independent and identically distributed from the all-to-all graph. An exponential multiplicative cooling schedule was used to cool to the final temperature over 5000 iterations.

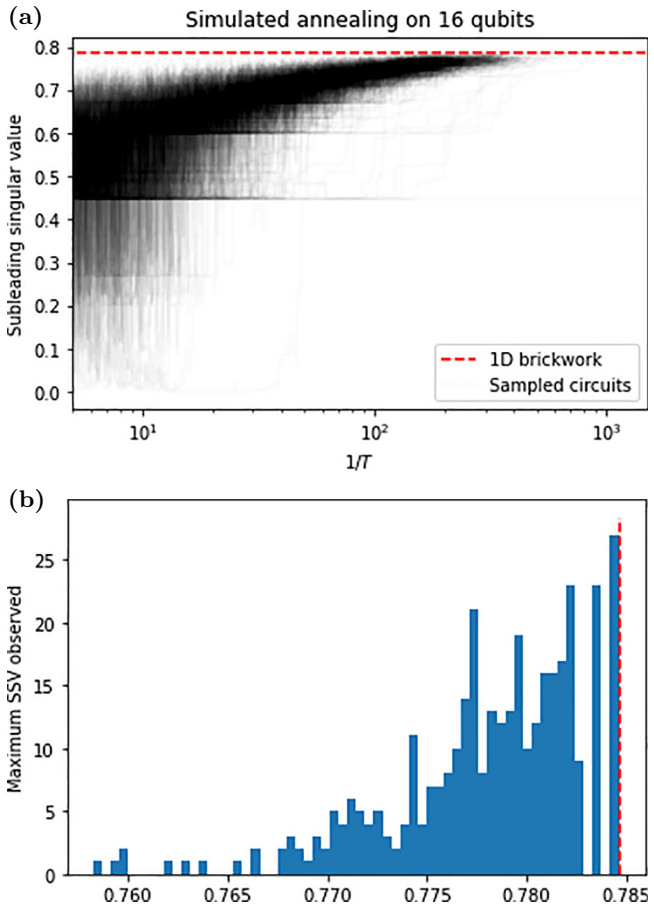

FIG. 12. (a) Subleading singular values by inverse temperature during 437 runs of the simulated annealing process. Here $N = 16$, $q = 2$, and $t = 2$. (b) Maximum SSVs attained by each run.

APPENDIX F: TIGHTER BOUNDS FOR BRICKWORK ARCHITECTURES OF ARBITRARY DIMENSION

For specific, well-connected architectures, we can use the cluster-merging picture to obtain a tighter bound than that of Theorem 1. One example is the generalized brickwork architectures on any dimension. We make the following claim.

Lemma 27. For $t = 2$, the effective connection depth ℓ on a generalized d-dimensional brickwork architecture on L^d sites is at most $2 + o_L(1)$. That is, the subleading singular value

$$s \leq 1 - (1 - s_{1D})^2 + o_L(1). \tag{F1}$$

In other words, we have taken the $2d$ layers of a d-dimensional brickwork's periodic block, and have effectively removed all but two of them from consideration in Theorem 7.

Each of the $2d$ layers in a brickwork architecture's periodic block chooses one of the d directions (horizontal, vertical, etc.) and one of two parities (odd or even). The cluster-merging picture of these layers comes in three stages. In the first stage, each layer has a different direction from all the layers below it. These layers combine finite [i.e., $O(1)$] size clusters into other finite size clusters. Each cluster after m such layers is an m-dimensional hypercube over 2^m sites. Depending on the order of the layers, this first stage can have any number of layers from 1 to d. The second stage begins once a layer has the same direction as a previous layer below it. In this stage, the hypercube

clusters get strung together along this direction, $L/2$ at a time, forming loops that can be reduced to periodic 1D brickwork architectures. This stage consists of only one layer.

After this stage, each cluster is of size $\Omega(L)$. Subsequent layers will connect these clusters either into pairs (if their same-direction counterpart was not applied yet) or $L/2$-length loops (if it was applied). Either way, each cluster will have $\Omega(L)$ connections with each neighbor in the graph. This effectively creates a 1D brickwork architecture with internal dimension $\tilde{q} = q^{\Omega(L)}$. Since the 1D brickwork singular value is $2\tilde{q}/(\tilde{q}^2+1) = O(\tilde{q}^{-1})$, all of these layers will have layer-restricted singular values that decay exponentially in L, and hence will contribute a vanishing amount $o_L(1)$ to Eq. (34).

We have a second stage with one layer that reduces to the periodic 1D brickwork architecture, and a third stage with vanishing contribution. Therefore, it remains to show that the combined subleading singular value of the first stage layers is bounded by the 1D brickwork architecture. In other words, we want to bound the subleading singular value of an arbitrary hypercube by the 1D brickwork architecture.

The subleading singular value of the 2D hypercube (i.e., a square of four sites) is $s_2 = 2q^2/(q^2+1)^2$. We can use the cluster-merging picture to bound hypercubes of higher dimension. At layer m, our clusters are size 2^{m-1}, and we are joining them together in pairs, with 2^{m-1} connections per pair. Since we have only two clusters in each connected section of our graph, finding the layer-restricted singular value $\mathfrak{s}_m$ just amounts to an optimization over four basis states. We have

$$\begin{aligned}\mathfrak{s}_m &= \frac{1}{(1-q^{-2^m})^2}\sqrt{2(1+q^{-2^m})\left[\frac{2}{q^2+1}\right]^{2^{m-1}} - 8q^{-2^m}} \\ &\le \sqrt{2}\left(\frac{16}{15}\right)\left(\frac{2}{q^2}\right)^{2^{m-2}} \\ &\le \sqrt{2}\left(\frac{16}{15}\right)\left(\frac{2}{q^2}\right)^{2(m-2)}, \end{aligned} \tag{F2}$$

where we bound the superexponential decay in m by an exponential for all $m \ge 3$. Then any $(d \ge 3)$-dimensional hypercube s_d has singular values bounded by

$$\begin{aligned} s_d^2 &\le 1 - (1-s_2^2)\prod_{m=3}^{d}(1-\mathfrak{s}_m^2) \\ &\le 1 - (1-s_2^2)\exp\left[-\sum_{m=3}^{d}\mathfrak{s}_m^2\right] \\ &\le 1 - (1-s_2^2)\exp\left[-2\left(\frac{16}{15}\right)^2\sum_{m=3}^{d}\left(\frac{2}{q^2}\right)^{4(m-2)}\right] \\ &\le 1 - \left(1-\frac{4q^4}{(q^4+1)^2}\right)\exp\left[-2\left(\frac{16}{15}\right)^2\frac{16}{q^8-16}\right]. \end{aligned} \tag{F3}$$

This decays more rapidly in q than the 1D brickwork architecture, and at $q = 2$ we get the bound $s_d(2) \le 0.478 < s_{1D}(2)$. Therefore, the hypercube singular value is below the 1D brickwork singular value for all $q \ge 2$.

A more thorough optimization over the size M of the first stage will probably get an overall singular value bound that is much closer to s_{1D}. This is because the size, and connectivity, of the clusters in the second stage is q^{2^M}, so it would probably produce a singular value that is close to $O(q^{-2^M})$. In other words, we are not allowed to make the first stage that deep, without reducing the layer-restricted singular value of the second stage.

[1] Joseph Emerson, Robert Alicki, and Karol Zyczkowski, Scalable noise estimation with random unitary operators, J. Opt. B: Quantum Semiclassical Opt. **7**, S347 (2005).

[2] Easwar Magesan, Jay M. Gambetta, and Joseph Emerson, Characterizing quantum gates via randomized benchmarking, Phys. Rev. A **85**, 042311 (2012).

[3] Patrick Hayden, Debbie Leung, Peter W. Shor, and Andreas Winter, Randomizing quantum states: Constructions and applications, Commun. Math. Phys. **250**, 371 (2003).

[4] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven, Characterizing quantum supremacy in near-term devices, Nat. Phys. **14**, 595 (2018).

[5] Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani, On the complexity and verification of quantum random circuit sampling, Nat. Phys. **15**, 159 (2019).

[6] Scott Aaronson and Lijie Chen, in *32nd Computational Complexity Conference, CCC 2017, July 6–9, 2017, Riga, Latvia*, edited by Ryan O'Donnell *LIPIcs* (Schloss Dagstuhl—Leibniz-Zentrum für Informatik, 2017), Vol. 79, p. 22:1.

[7] Adam Bouland, Bill Fefferman, Zeph Landau, and Yunchao Liu, in *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA*, pp. 1308–1317.

[8] Ramis Movassagh, The hardness of random quantum circuits, Nat. Phys. **19**, 1719 (2023).

[9] Yasuhiro Kondo, Ryuhei Mori, and Ramis Movassagh, in *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA*, pp. 1296–1307.

[10] Patrick Hayden and John Preskill, Black holes as mirrors: Quantum information in random subsystems, J. High Energy Phys. **2007**, 120 (2007).

[11] Yasuhiro Sekino and L. Susskind, Fast scramblers, Journal of High Energy Physics **2008**, 065 (2008).

[12] Patrick Hayden, Sepehr Nezami, Xiao Liang Qi, Nathaniel Thomas, Michael Walter, and Zhao Yang, Holographic duality from random tensor networks, J. High Energy Phys. **2016**, 9 (2016).

[13] Brian Skinner, Jonathan Ruhman, and Adam Nahum, Measurement-induced phase transitions in the dynamics of entanglement, Phys. Rev. X **9**, 031009 (2019).

[14] Yunxiang Liao and Victor Galitski, Effective field theory of random quantum circuits, Entropy **24**, 823 (2022).

[15] Frank Arute, *et al.*, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).

[16] A. Pirker, J. Wallnöfer, and W. Dür, Modular architectures for quantum networks, New J. Phys. **20**, 53054 (2018).

[17] G. Vidal, Class of quantum many-body states that can be efficiently simulated, Phys. Rev. Lett. **101**, 110501 (2008).

[18] Isaac H. Kim and Brian Swingle, Robust entanglement renormalization on a noisy quantum computer, ArXiv:1711.07500.

[19] Joseph Emerson, Etera Livine, and Seth Lloyd, Convergence conditions for random quantum circuits, Phys. Rev. A—At., Mol. Opt. Phys. **72**, 060302 (2005).

[20] Andris Ambainis and Joseph Emerson, *22nd Annual IEEE Conference on Computational Complexity* (IEEE Computer Society, San Diego, CA, 2007), pp. 129–140.

[21] Alexander M. Dalzell, Nicholas Hunter-Jones, and Fernando G. S. L. Brandão, Random quantum circuits anticoncentrate in log depth, PRX Quantum **3**, 10333 (2022).

[22] Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki, Local random quantum circuits are approximate polynomial-designs, Commun. Math. Phys. **346**, 397 (2016).

[23] Jonas Haferkamp, Random quantum circuits are approximate unitary t-designs in depth $O(nt^{5+o(1)})$, Quantum **6**, 795 (2022).

[24] Nicholas Hunter-Jones, Unitary designs from statistical mechanics in random quantum circuits, ArXiv:1905.12053.

[25] Aram W. Harrow and Saeed Mehraban, Approximate unitary t-designs by short random quantum circuits using nearest-neighbor and long-range gates, Commun. Math. Phys. **401**, 1531 (2023).

[26] Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang, Random unitaries in extremely low depth, ArXiv:2407.07754.

[27] Nicholas LaRacuente and Felix Leditzky, Approximate unitary k-designs from shallow, low-communication circuits, ArXiv:2407.07876.

[28] Michał Oszmaniec, Michał Horodecki, and Nicholas Hunter-Jones, Saturation and recurrence of quantum complexity in random quantum circuits, ArXiv:2205.09734.

[29] Shivan Mittal and Nicholas Hunter-Jones, Local random quantum circuits form approximate designs on arbitrary architectures, ArXiv:2310.19355.

[30] Chi-Fang Chen, Jeongwan Haah, Jonas Haferkamp, Yunchao Liu, Tony Metger, and Xinyu Tan, Incompressibility and spectral gaps of random circuits, ArXiv:2406.07478.

[31] See Definition 2 or Fig. 1.

[32] For architectures without any regular structure, $\bar{\ell}$ may itself depend on depth, so this is not always an explicit formula for the t-design depth. See Theorem 12 for a precise definition of $\bar{\ell}$.

[33] Winton G. Brown and Lorenza Viola, Convergence rates for arbitrary statistical moments of random quantum circuits, Phys. Rev. Lett. **104**, 250501 (2010).

[34] Note that our bound is based on a fixed 1D brickwork instead of a random layer ordering. This allows us to tighten the bound of Ref. [22] for parallel local random circuits by a factor of 2, by skipping the factor of $\frac{1}{2}$ relaxation of λ_2 in their Lemma 21.

[35] There is a factor-of-2 difference between our definition of C and that of Ref. [24]. This is because we are counting periods instead of layers.

[36] There can be no single domain wall in periodic boundary conditions, so the leading term is now given by the two-domain-wall sector.

[37] The connectedness is relevant because it allows the use of SWAP gates to obtain any pairwise operation.

[38] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. Divincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter, Elementary gates for quantum computation, Phys. Rev. A **52**, 3457 (1995).

[39] We can further tighten the bound for certain architectures by observing that ℓ_i only needs to count layers that merge clusters. Layers consisting of only internal edges within already-merged clusters need not be counted.

[40] Aram W. Harrow and Richard A. Low, Random quantum circuits are approximate 2-designs, Commun. Math. Phys. **291**, 257 (2009).

[41] Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh Vazirani, Quantum Hamiltonian complexity and the detectability lemma, ArXiv:1011.3445.

[42] The upper bound is $N-1$ because we ignore layers with completely disconnected cluster-merging graphs.

[43] Paul Erdős and Alfréd Rényi, On random graphs I, Publ. Math. Debrecen **6**, 18 (1959).

[44] Paul Erdős and Alfréd Rényi, On a classical problem of probability theory, Magyar Tud. Akad. Mat. Kutató Int. Közl. **6**, 215 (1961).