Random Quantum Circuits Anticoncentrate in Log Depth

Alexander M. Dalzell[◉],[1] Nicholas Hunter-Jones[◉],[2] and Fernando G. S. L. Brandão[1,3]

[1]*Institute for Quantum Information and Matter, Caltech, Pasadena, California 91125, USA*
[2]*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*
[3]*AWS Center for Quantum Computing, Pasadena, California 91125, USA*

We consider quantum circuits consisting of randomly chosen two-local gates and study the number of gates needed for the distribution over measurement outcomes for typical circuit instances to be anticoncentrated, roughly meaning that the probability mass is not too concentrated on a small number of measurement outcomes. An understanding of the conditions for anticoncentration is important for determining which quantum circuits are difficult to simulate classically, as anticoncentration has been in some cases an ingredient of mathematical arguments that simulation is hard and in other cases a necessary condition for easy simulation. Our definition of anticoncentration is that the expected *collision probability* of the distribution—that is, the probability that two independently drawn outcomes will agree—is only a constant factor larger than the collision probability for the uniform distribution. We show that when the two-local gates are each drawn from the Haar measure (or any 2-design), at least $\Omega(n \log(n))$ gates (and thus $\Omega(\log(n))$ circuit depth) are needed for this condition to be met on an n-qudit circuit. In both the case where the gates are nearest neighbor on a one-dimensional ring and the case where gates are long range, we show that $O(n \log(n))$ gates are also sufficient and we precisely compute the optimal constant prefactor for the $n \log(n)$. The technique we employ relies upon a mapping from the expected collision probability to the partition function of an Ising-like classical statistical-mechanical model, which we manage to bound using stochastic and combinatorial techniques.

I. INTRODUCTION

Random quantum circuits (RQCs) are a crucial model for understanding a diverse set of phenomena in both quantum information and quantum many-body physics. They have been used to study the onset of quantum chaos and the dynamical spread of entanglement in strongly interacting quantum systems [1–3], including information processing in black holes [4]. They also form the basis for recent experiments aiming to demonstrate exponential quantum advantage [5–7].

The utility of RQCs in these situations derives from a myriad of quantitative properties that they have been shown to possess. For example, RQCs quickly generate entanglement [1,8,9], lead to fast scrambling and decoupling of quantum information [10,11], and act as efficient encoding circuits for good quantum error-correcting codes [12]. When the circuits are geometrically local, they

lead to ballistic spreading of local operators [2,3]. Furthermore, they form approximate unitary designs, that is, despite being composed of local gates, they efficiently approximate a global random unitary transformation up to any polynomial number of moments [13–17]. Meanwhile, the computing of transition amplitudes of RQCs has been shown to be just as difficult as for arbitrary quantum circuits [18–22], suggesting that classical simulation of RQCs should require exponential time.

In this work, we focus on another property of RQCs called *anticoncentration*. Roughly speaking, when we measure the output state of the circuit in the computational basis, anticoncentration is the property that the distribution over measurement outcomes is fairly well spread across all possible outcomes and not too concentrated onto just one or only a small portion of those outcomes. Quantitatively, our definition of anticoncentration depends on the *collision probability*, the probability that measurement outcomes from two independent copies of the circuit agree. A RQC architecture is said to be anticoncentrated if the collision probability is at most a constant factor larger than its minimal value. An understanding of when this is the case is particularly important for knowing when RQCs are hard to classically simulate. On the one hand,

anticoncentration is a necessary ingredient in most formal hardness arguments for RQC simulation [18,23–29]. On the other hand, certain classical algorithms for simulating RQCs require anticoncentration in order to be efficient; for example, the algorithms discussed in Refs. [30,31] for noisy-circuit simulation and the algorithm in Ref. [32] that spoofs the linear cross-entropy benchmarking metric introduced in Ref. [6].

In most previous work where RQC anticoncentration is needed, it has been asserted as an implication of the 2-design property (see, e.g., Refs. [27,33]). However, the 2-design property is much stronger than is required for anticoncentration. It has been shown that n-qubit RQCs on a fully connected architecture form approximate 2-designs after roughly $O(n)$ depth [13] and this has later been shown to also apply to geometrically local RQCs in one dimension and improved to $O(n^{1/D})$ in D spatial dimensions [16]. However, recent work by Barak, Chou, and Gao [32]—using a similar method to the one presented here—has shown that for one-dimensional (1D) RQCs, the collision probability converges in depth $O(\log(n))$, much faster than the 2-design depth of $O(n)$. They have also conjectured that two-dimensional (2D) RQCs anticoncentrate in depth $O(\sqrt{\log(n)})$.

In this work, we prove sharp bounds on the number of gates needed for anticoncentration in two RQC architectures. For 1D RQCs, we confirm the $O(\log(n))$ upper bound on the anticoncentration depth in Ref. [32] and add a lower bound that matches the upper bound even up to the constant prefactor of the $\log(n)$. We also show that an $\Omega(\log(n))$ lower bound on the depth needed for anticoncentration holds regardless of which RQC architecture we use, which refutes the conjecture from Ref. [32] that 2D RQCs anticoncentrate in $O(\sqrt{\log(n)})$ depth. We then consider a fully connected (i.e., not geometrically local) RQC architecture, where each gate acts on a pair of qudits chosen randomly among all $n(n-1)/2$ possible such pairs. We show that, for qubits (local dimension $q = 2$), $5n\log(n)/6$ gates are necessary and sufficient (up to subleading corrections) for anticoncentration to be achieved, which settles a conjecture in Ref. [16].

Our method employs a technique for analyzing RQCs that converts the collision probability into a weighted sum over bit assignments to each location in the circuit diagram; this weighted sum can be viewed as a partition function for an Ising-like statistical-mechanical model. The bit assignments can also be interpreted as a Markov chain and the number of gates needed for anticoncentration ultimately translates into the time needed for certain expectation values to converge under the dynamics of the Markov chain. This method not only yields sharp quantitative bounds but it also produces an appealing qualitative explanation on how and why the collision probability reaches its limiting value, which allows for effective

heuristic reasoning even in architectures that we do not explicitly consider here.

The main takeaways from our work are twofold. First, we show that anticoncentration is generally achieved much faster than the 2-design property. The fact that anticoncentration occurs in $\Theta(n\log(n))$ circuit size both in 1D and for the fully connected architecture—these being two opposite extremes of geometric locality—suggests that anticoncentration may require only $\Theta(n\log(n))$ size for *any* reasonably well-connected architecture. This comes in sharp contrast to the situation for unitary designs, where the scaling of the size needed with n is highly dependent on the architecture. Second, the fact that we can prove tight upper and lower bounds suggests a broader utility for our method based on the correspondence between RQCs and statistical-mechanical partition functions.

II. ANTICONCENTRATION AND THE COLLISION PROBABILITY

A RQC architecture is an instruction set on how to draw a circuit diagram given the number of qudits n (each with local Hilbert space dimension q) and the size s of the circuit. The two architectures we consider specifically are the *1D architecture* (with periodic boundary conditions) and the *complete-graph* architecture. The associated RQC ensemble for a RQC architecture is formed by following this instruction set and then choosing the value of each gate in the diagram independently and uniformly at random from the Haar measure. If we fix an instance U from this ensemble, there is an associated output probability distribution p_U over q^n possible computational basis measurement outcomes $x \in [q]^n$ (where $[q] = \{1, 2, \ldots, q\}$). Anticoncentration tries to capture the notion that the probability mass is well spread out over all the outcomes. The uniform distribution, where each output is allocated q^{-n} fraction of the total probability mass, is the ultimate anticoncentrated distribution because the mass is exactly equally spread but we say that a distribution is still anticoncentrated as long as the average fluctuations from uniform are no larger than $O(q^{-n})$. This definition is captured precisely by the *collision probability*, which is $\sum_x p_U(x)^2$. The collision probability gives the probability that measurement outcomes from two independent copies of the circuit are identical. It is also proportional to the second moment (and thus is related to the variance) of the output probability of a randomly chosen bit string. If p_U is the uniform distribution, then the collision probability is q^{-n}, its minimal possible value. For a RQC architecture at a specified qubit number n and circuit size s, we consider the collision probability averaged over the randomly chosen circuit instances U:

$$Z := \mathbb{E}_U \left[\sum_{x \in [q]^n} p_U(x)^2 \right] = q^n \mathbb{E}_U \left[p_U(1^n)^2 \right], \qquad (1)$$

where the second equality holds because by symmetry each of the q^n terms in the sum yields the same number under expectation as long as at least one Haar-random gate acts on each qudit.

We say that a RQC architecture with n qudits and s gates is anticoncentrated if there is a constant α (independent of n) with $0 < \alpha \leq 1$ for which $Z \leq \alpha^{-1} q^{-n}$, i.e., that the collision probability is only a constant factor larger than its minimal value. In particular, our theorem statements roughly correspond to the choice $\alpha = 1/4$ but other choices of α would yield the same results up to leading order. If desired, Markov's inequality can then be used to bound the fraction of the randomly chosen U the collision probability of which is larger than some constant multiple of Z. Moving forward, for convenience, when we say collision probability we mean the average collision probability Z.

Very shallow circuit architectures are not anticoncentrated: there are expected to be some output probabilities x for which $p_U(x)$ is exponentially larger than the mean of q^{-n}. As the circuit gets deeper, we expect the probability distribution to become closer to uniform, but even at infinite depth, when the circuit unitary U becomes a globally Haar-random $q^n \times q^n$ unitary, the output distribution still does not become completely uniform. In this case, the output distribution will typically follow a Porter-Thomas distribution [34] and Z can be exactly computed as

$$\lim_{s \to \infty} Z = Z_H := \frac{2}{q^n + 1}, \qquad (2)$$

roughly twice as large as the minimal value of q^{-n} associated with the uniform distribution. This statement is proved using the techniques described later. For a graphical illustration of these cases, see Fig. 1.

While one could capture the notion of anticoncentration with a different definition, the definition we choose is useful and relevant because it has concrete ramifications in all of the previously mentioned applications of anticoncentration. For example, one implication of our definition (by application of the Paley-Zygmund inequality) is that if $Z \leq \alpha^{-1} q^{-n}$, then for any $0 \leq \beta \leq 1$,

$$\Pr_U[p_U(x) \geq \beta q^{-n}] \geq \alpha (1 - \beta)^2, \qquad (3)$$

meaning that for at least a constant fraction of the circuit instances, the probability of a given measurement outcome x is at least a constant multiple β of the mean measurement probability q^{-n}. This sort of inequality is the relevant one for turning good additive approximations into good multiplicative approximations (with reasonable probability), employed in, e.g., Refs. [18,23–29] to argue [35] that it is hard to classically sample output distributions for a large fraction of instances up to small total variation distance error (for more details, see Sec. B). In fact, equations such as Eq. (3) are sometimes taken to be the definition of anticoncentration [27], which is a weaker definition than ours since, in principle, Eq. (3) can hold even in cases where Z exceeds any constant multiple of q^{-n}.

III. OUR RESULTS

We show that the collision probability is given by a discrete sum, which we interpret as the expectation value of a certain stochastic process. The correspondence
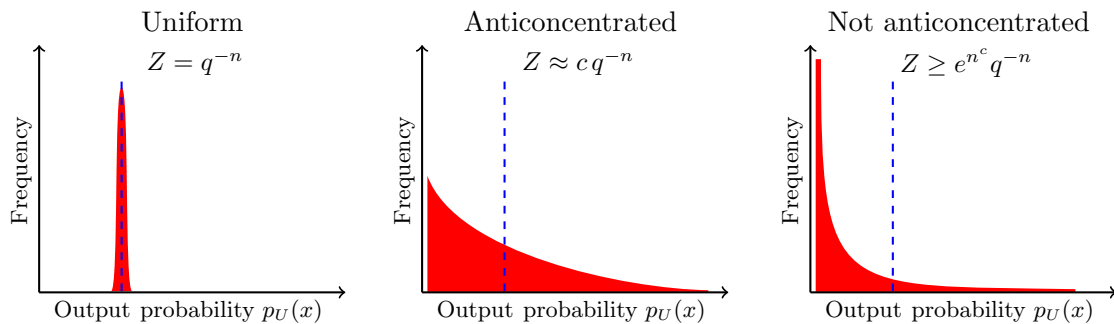

FIG. 1. A sketch of anticoncentration. In the three examples, the fraction (frequency) of bit strings x for which $p_U(x) = p$ is plotted against p. (This could be either for a fixed choice of U or averaged over random choice of U.) Since there are q^n bit strings, the mean of this distribution is q^{-n} (dotted blue line). For the uniform distribution, which is completely anticoncentrated, all q^n outcomes are allocated probability mass q^{-n} and the collision probability is $Z = q^{-n}$. For globally Haar-random unitaries, the output probabilities are on average q^{-n} but have some nonzero variance and the collision probability is $Z \approx 2q^{-n}$. Whenever $Z \approx cq^{-n}$ for some c independent of n, we call the distribution anticoncentrated. For low-depth RQCs, the mean output probability is q^{-n} but the variance is much larger and the collision probability is much larger than q^{-n}. Most of the probability mass is concentrated onto a few measurement outcomes, while the vast majority of the outcomes are assigned a very small amount of mass, leading to the divergence in the frequency that $p_U(x)$ is close to 0 depicted in the plot.

between the collision probability and the discrete sum is described in Sec. V and a complete derivation is provided in Appendix B.

In analyzing our expression, we derive rigorous upper and lower bounds on the collision probability generally and for two specific architectures. These bounds are stated here and the proofs are provided in the appendices. These bounds are then used to form upper and lower bounds quoted in Table I on the anticoncentration size s_{AC}, defined as the minimum circuit size required such that $Z \leq 2Z_H$. The constant 2 in the definition of s_{AC} is arbitrary but, since we show that Z approaches Z_H as $Z_H[1 + e^{-\Omega(s/n)}]$, a different choice of constant would only lead to linear-in-n changes to s_{AC}, which would be subleading and would not affect any of the statements in Table I. All logarithms in this paper are natural logarithms.

A. Collision probability upper bounds

Our upper bounds take the following form:

$$Z \leq Z_H \left[1 + e^{-\frac{2a}{n}(s-s^*)} \right], \tag{4}$$

where the constant a is independent of n and depends on the circuit architecture and s^* is a function of n that also depends on the architecture. Thus, if the anticoncentration size s_{AC} is defined to be the minimum size s such that $Z \leq 2Z_H$, then we have $s_{AC} \leq s^*$. Specifically, we have the following results, which are restated here as theorems and proved rigorously in the appendices.

First, we consider the *1D architecture* with periodic boundary conditions, where the qudits are arranged on a ring and alternating layers of $n/2$ nearest-neighbor Haar-random gates are applied.

Theorem 1. *For the 1D architecture, Eq. (4) holds with*

$$a = \log\left(\frac{q^2+1}{2q} \right), \tag{5}$$

$$s^* = \frac{1}{2a} n \log(n) + n \left[\frac{1}{2a} \log(e-1) + \frac{1}{2} \right], \tag{6}$$

whenever $s \geq s^*$.

Since this depth of the 1D architecture is given by $d := 2s/n$, we can define $d^* := 2s^*/n = a^{-1}\log(n) + O(1)$ for 1D and conclude that the "anticoncentration depth" d_{AC} satisfies $d_{AC} \leq d^* = O(\log(n))$.

Similarly, we show an upper bound for the *complete-graph architecture*, where each gate acts on a random pair of qudits without regard for their spatial proximity.

Theorem 2. *For the complete-graph architecture, Eq. (4) holds with*

$$a = \frac{(q-1)^2}{2(q^2+1)}, \tag{7}$$

$$s^* = \frac{q^2+1}{2(q^2-1)} n \log(n) + cn, \tag{8}$$

whenever $s \geq s^*$, for a constant c that is independent of n.

A size-s circuit diagram chosen randomly from the complete-graph architecture will have depth at most $O(s\log(n)/n)$ with high probability [11], meaning that $O(\log(n)^2)$ depth is typically sufficient for anticoncentration in the complete-graph architecture.

We also consider general architectures. We define a property called *regularly connected* (Definition 5 in Appendix C), which applies to a RQC architecture when for any partition of qubits into two sets, there will be a gate in the circuit that couples the two sets at least once every $O(n)$ gates. Nearly all natural architectures have this property, including standard architectures in D spatial dimensions for any D.

Theorem 3. *If an architecture is regularly connected, then Eq. (4) holds with $a = \Theta(1)$ and $s^* = \Theta(n^2)$*

This corresponds to $\Theta(n)$ gates per qudit. This result is weaker than our specific result for the 1D and complete-graph architecture and we conjecture that much better is possible.

Conjecture 1. *Theorem 3 can be improved to $s^* = \Theta(n\log(n))$.*

TABLE I. Summary of results: upper and lower bounds on the circuit size s_{AC} at which anticoncentration is achieved for different random circuit architectures.

Architecture	s_{AC} upper bound	s_{AC} lower bound
General	$O(n^2)$	$\Omega(n\log(n))$
1D	$\left[2\log\left(\frac{q^2+1}{2q} \right) \right]^{-1} n\log(n) + O(n)$	$\left[2\log\left(\frac{q^2+1}{2q} \right) \right]^{-1} n\log(n) - O(n)$
Complete graph	$\frac{q^2+1}{2(q^2-1)} n\log(n) + O(n)$	$\frac{q^2+1}{2(q^2-1)} n\log(n) - O(n)$

B. Collision probability lower bounds

Our lower bounds on the collision probability take the form

$$Z \geq \frac{Z_H}{2} \exp\left[A e^{\log(n) - Bs/n}\right], \qquad (9)$$

for constants A and B that are independent of n. (The lower bound for the complete-graph architecture takes a different but very similar form.) This form implies that if s grows with n like $s \approx fn \log(n)/B$ for some $f < 1$, then we have $Z/Z_H \geq (1/2)e^{An^{1-f}}$, which becomes arbitrarily large as $n \to \infty$, meaning that the architecture is not anticoncentrated. This puts a lower bound on the "anticoncentration" size s_{AC} of $s_{\mathrm{AC}} \geq n \log(n)/B - O(n)$.

Specifically, we show a general lower bound, as well as specific lower bounds for the 1D and complete-graph architectures.

Theorem 4. *For any RQC architecture with s two-qudit gates, the following holds:*

$$Z \geq \frac{Z_H}{2} \exp\left[\frac{\log(q)}{q+1} e^{\log(n) - 2\log(q^2+1)s/n}\right]. \qquad (10)$$

This has the consequence that if s_{AC} and d_{AC} are defined as the minimum size and minimum depth for which $Z \leq 2Z_H$, then

$$s_{\mathrm{AC}} \geq [2\log(q^2+1)]^{-1} n \log(n) - O(n), \qquad (11)$$

$$d_{\mathrm{AC}} \geq [\log(q^2+1)]^{-1} \log(n) - O(1). \qquad (12)$$

We improve on the general lower bound for the two specific architectures that we consider.

Theorem 5. *For the 1D architecture, there exists a constant A such that*

$$Z \geq \frac{Z_H}{2} \exp\left[A e^{\log(n) - 2as/n}\right], \qquad (13)$$

where $a = \log[(q^2+1)/(2q)]$ is the same as for the upper bound in Eq. (5).

This implies that in 1D,

$$s_{\mathrm{AC}} \geq (2a)^{-1} n \log(n) - O(n), \qquad (14)$$

$$d_{\mathrm{AC}} \geq a^{-1} \log(n) - O(1), \qquad (15)$$

which is tight with the upper bound up to subleading corrections.

Theorem 6. *For the complete-graph architecture,*

$$Z \geq \frac{Z_H}{2} \exp\left\{\frac{\log(q)}{q+1} e^{\log(n) + \log\left[1 - \frac{2(q^2-1)}{n(q^2+1)}\right]s}\right\}. \qquad (16)$$

Although a slightly different form than the other lower bounds, this still yields the conclusion

$$s_{\mathrm{AC}} \geq \frac{q^2+1}{2(q^2-1)} n \log(n) - O(n), \qquad (17)$$

which is tight with the upper bound up to subleading corrections. When $q = 2$ (qubits), the prefactor of the $n \log(n)$ is $5/6$, settling a conjecture proposed in Ref. [16].

The upper and lower bounds together allow us to conclude that $s_{\mathrm{AC}} = \Theta(n \log(n))$ for both the 1D architecture and the complete-graph architecture and, in fact, we have matching upper and lower bounds on the constant prefactor of the $n \log(n)$.

We note that for $q \geq 5$, our results have the counterintuitive implication that the 1D architecture anticoncentrates faster than the complete-graph architecture, even though it is geometrically local. We argue that this is an artifact of the definition of the models and can be explained by the fact that the qudit pairs acted upon by the gates in the complete-graph architecture are chosen randomly, while the qudit pairs in the 1D architecture are not random; in fact, in the latter case they are optimally packed into layers of $n/2$ nonoverlapping gates. As q increases, anticoncentration becomes arbitrarily fast for the 1D architecture [the coefficient of the $n \log(n)$ decreases like $1/\log(q)$]. Meanwhile, for the complete-graph architecture, no matter how large q is, there will always be some minimum number of gates—roughly $n \log(n)/2$—needed simply to guarantee that all the qudits have been involved in the circuit with high probability. We suspect that a parallelized version of the complete-graph architecture would anticoncentrate with a slightly better constant than the 1D architecture.

IV. RELATED WORK AND IMPLICATIONS

Here, we highlight a few relevant previous works and emphasize how our results fit in.

(a) Harrow and Mehraban [16] have studied how quickly RQCs form approximate unitary t-designs and anticoncentrate for various architectures. For geometrically local circuits, they have shown that the approximate t-design property is achieved after only $O(n^{1/D})$ depth in D spatial dimensions, the first work to break the $O(n)$ barrier for designs. Since anticoncentration follows from the approximate 2-design property, their work implies an $O(n^{1/D})$ upper bound on the anticoncentration depth. We show that for $D = 1$, the anticoncentration depth is actually $\Theta(\log(n))$ and we conjecture that this is also the case for $D \geq 2$ but we do not prove this, so the $O(n^{1/D})$ bound remains the best known for $D \geq 2$. Harrow and Mehraban have also considered the question of anticoncentration in the complete-graph

architecture and have shown an upper bound on the anticoncentration size of $O(n \log(n)^2)$ and a lower bound of $\Omega(n \log(n))$. They have used heuristic reasoning to conjecture that (for $q = 2$) the anticoncentration size should be $5n \log(n)/3$, up to leading order. We are able to show that to leading order, the anticoncentration size for the complete-graph architecture is $5n \log(n)/6$. This is off by a factor of 2 from the conjecture stated in their paper, which we suspect is due to a minor error in their heuristic reasoning.

(b) Barak, Chou, and Gao [32] have developed a classical algorithm for shallow RQCs that achieves a non-negligible score on the linear cross-entropy benchmarking (XEB) metric despite not performing a full simulation of the RQCs. The linear XEB metric has been used by Google to verify its 2019 quantum computational supremacy experiment [6]. Barak, Chou, and Gao show that if a depth-d RQC architecture in D spatial dimensions has collision probability Z, their algorithm achieves a score of ϵ with high probability after a total runtime $(2^n Z) \exp(\epsilon 15^{-d}) \mathrm{poly}(n, 2^{d^D})$ (here, $q = 2$). They prove that $Z = O(2^{-n})$ after $d = \Omega(\log(n))$ for 1D RQCs, which is equivalent to our Theorem 1. This shows that their algorithm achieves an $\epsilon \geq 1/\mathrm{poly}(n)$ score in polynomial time for logarithmic depth 1D RQCs. For 2D RQCs, they conjecture that $Z = O(2^{-n})$ after depth $d = O(\sqrt{\log(n)})$, which would imply that their algorithm achieves an $\epsilon \geq 1/\mathrm{poly}(n)$ score in polynomial time at that depth. Our Theorem 4 contradicts their conjecture by showing generally that $2^n Z \geq \exp[n^{1-o(1)}]$ when d is sublogarithmic.

(c) Our method performs expectations over individual gates in the RQC using formulas for Haar integration, a strategy that has also been used on similar problems in the past. Many works have used this strategy to form a random walk over Pauli strings with wide-ranging applications [8,10–13,16,36–38]. Our analysis applies this strategy in a distinct way that more closely resembles a series of works that interpret the resulting expression as the partition function of classical statistical-mechanical models [2,3,17,39–44]. Here, we analyze those partition functions using a Markov-chain analysis but our Markov chain has different transition rules compared to the Pauli-string Markov chain.

A. Connection to 2-design

Anticoncentration for RQCs (as well as some Hamiltonian models) is often established as a consequence of the convergence to approximate unitary 2-designs, where approximately reproducing the first two moments of the

Haar measure allows one to bound the RQC collision probability. For both 1D and complete-graph RQCs, size-$O(n^2)$ circuits (of linear depth) form approximate 2-designs and therefore anticoncentrate. There are a number of definitions of approximate unitary designs utilizing different norms; we briefly comment on the definitions and requirements for anticoncentration in this architecture.

As we review in Appendix F, defined in terms of the diamond norm, ε-approximate 2-designs have a collision probability upper bounded by Z_H up to additive error. In order to achieve anticoncentration, ε must be taken to be exponentially small (i.e., we require $\varepsilon = 1/q^{2n}$). A stronger notion of approximate design in terms of the complete positivity of the difference in channels has been introduced in Ref. [15]. Under this strong definition, 2-designs bound the collision probability up to relative error with respect to the Haar value and thus anticoncentrate. A much weaker definition of approximate design is the operator norm of the moment operators, often called the tensor-product-expander (TPE) condition. Interestingly, TPEs also bound the collision probability up to additive error but, again, the error needs to be exponentially small to achieve anticoncentration.

RQCs on the 1D architecture form ε-approximate 2-designs, in both diamond norm and the stronger definition, when the circuit size is $O(n(n + \log(1/\varepsilon)))$. Moreover, 1D random circuits actually form ε-approximate TPEs in constant depth, when the circuit size is $O(n \log(1/\varepsilon))$. But again, anticoncentration requires that ε be taken to be $\varepsilon = 1/q^{2n}$, thus mandating linear depth. So to establish that the collision probability is bounded up to a relative error, as in the definition of anticoncentration, the use of unitary 2-designs or a general bound on the moments necessitates linear depth. For nonlocal RQCs defined on a complete-graph, the best known upper bounds on the approximate 2-design depth are the same as for the 1D architecture. But it has been conjectured that this may be improved for nonlocal RQCs, which would close the gap between the 2-design time and the depth required for anticoncentration.

To further emphasize the distinction between anticoncentration and unitary 2-designs, we note that anticoncentration can be achieved for specific short-depth circuits without generating entanglement across the system (indeed, a circuit consisting of a single layer of single-qubit Hadamards suffices). Moreover, for an ensemble of RQCs, anticoncentration can be equivalently phrased as the statement that certain matrix elements of second-moment operator $\mathbb{E}_U[U^{\otimes 2} \otimes U^{*\otimes 2}]$ reach the Haar value of $2/q^{2n}$ after some depth, whereas the approximate 2-design condition gives that $\mathbb{E}_U[\langle \psi | U^{\otimes 2} \otimes U^{*\otimes 2} | \psi \rangle]$ is small for all states $|\psi\rangle$, even those that are entangled across the tensor copies. As we show in Appendix F, there are necessarily some states that require linear depth to equilibrate to the minimal Haar value, at least for RQCs on the 1D architecture.

B. Implications for arguments on hardness of simulation

Anticoncentration is a key ingredient in hardness-of-simulation arguments [18–29] that underlie quantum computational supremacy proposals. In this section, we roughly explain its role in those arguments and the implications our results have in this context.

The starting point for these hardness arguments is the long-known observation that the answer to a hard classical problem can be encoded into the output probability $p_U(x)$ of a quantum circuit U [45]. Thus, exact computation of $p_U(x)$ for arbitrary U and x should not be possible in classical polynomial time. This remains true even if one only needs to compute $p_U(x)$ up to some constant relative error. The ultimate goal in the context of quantum computational supremacy is to show that there is no polynomial-time classical algorithm that approximately simulates random circuits (or, at least, to give extremely convincing evidence in favor of this conclusion). More precisely, the approximate simulation task is to produce samples from a distribution p'_U for which

$$\|p_U - p'_U\|_1 := \sum_x |p_U(x) - p'_U(x)| = \varepsilon \quad (18)$$

for some small $\varepsilon = O(1)$ and to do this for a large fraction of U drawn randomly from some random ensemble. Turning the starting point into the ultimate goal requires a few steps (some of which rely on conjecture). Anticoncentration is one of these steps.

The primary role that anticoncentration plays is to turn a small *additive* difference $|p_U(x) - p'_U(x)|$ for most x into a small *relative* difference $r(x)$ for most x, where

$$r(x) := \frac{|p_U(x) - p'_U(x)|}{p_U(x)}. \quad (19)$$

If Eq. (18) is obeyed, then the value of $|p_U(x) - p'_U(x)|$ is on the order of ε/q^n for most x. Meanwhile, the mean value of $p_U(x)$ for random x is exactly $1/q^n$. If $p_U(x)$ is anticoncentrated, then for most x, $p_U(x)$ will be within a constant factor of the mean, as shown in the middle diagram of Fig. 1, and $r(x) = O(\epsilon)$ will hold for most x. However, if $p_U(x)$ is not anticoncentrated, then $p_U(x)$ will be much smaller than the mean for most x, as depicted in the right-hand diagram of Fig. 1. This means that without anticoncentration, $r(x) \gg \epsilon$ for most x, which is problematic because the hard classical problems encoded into $p_U(x)$ are no longer hard when the relative error is extremely large, so anticoncentration appears to be necessary if there is any hope of completing the hardness argument using existing techniques.

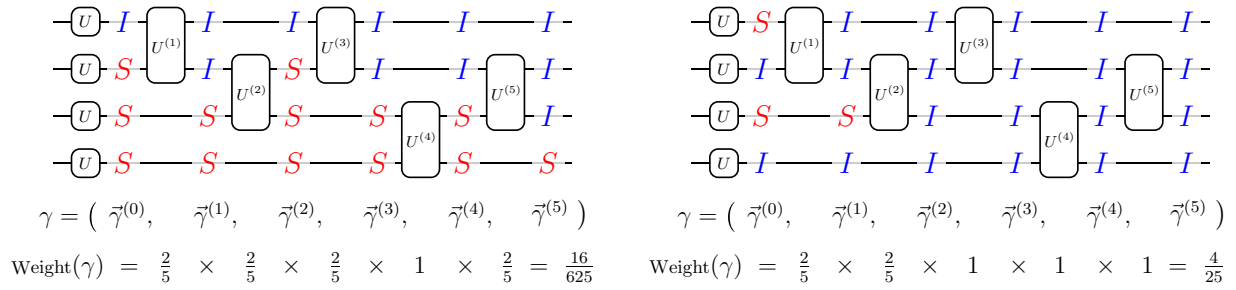Even if anticoncentration holds, more is needed to show hardness of approximately simulating RQCs. One must turn hardness of computing $p_U(x)$ into hardness of *sampling* from p_U and also turn hardness for arbitrary U into hardness for a *random U*. There are techniques that work for each of these steps individually but currently they do not work together simultaneously and thus an additional conjecture must be made.

Our work puts sharp bounds on the number of gates needed for anticoncentration to hold in multiple RQC architectures, which constrains when these hardness arguments have the potential to work. Our finding that the number of gates per qudit needed for anticoncentration grows only like $O(\log(n))$ in the 1D and complete-graph architectures implies that perhaps RQC-based quantum computational supremacy could be achieved at a shallower circuit depth than previously believed. For example, Google's 2019 quantum computational supremacy experiment was based on 2D RQCs of depth exceeding the $\sqrt{n}$ diameter of the qubit array [5,6]. The fact that 1D circuits anticoncentrate in $\Theta(\log(n))$ depth is evidence that 2D circuits should have the same scaling (if anything, anticoncentration should happen faster in 2D). Thus a similar quantum computational supremacy experiment might be equally defensible at $\Theta(\log(n))$ depth instead of $\Omega(\sqrt{n})$ depth. We note, however, that there are other reasons to want to go to larger depth (e.g., classical simulation via tensor-network methods becomes harder at larger depths).

Without anticoncentration, the hardness-of-simulation arguments appear to break down but this does not generally imply that simulation is easy. On this topic, a subset of these authors and others have described an algorithm for solving the approximate simulation problem for 2D RQCs [43]. The algorithm is proved to be efficient for a certain constant-depth (and thus not anticoncentrated) 2D RQC architecture but it is conjectured to become inefficient once the depth exceeds a larger constant threshold. Thus, the complexity of the algorithm transitions to inefficient before the circuits become anticoncentrated, suggesting that in 2D there could be a regime where the RQCs are too shallow to be anticoncentrated but classical simulation is still hard.

V. COLLISION PROBABILITY AS A SUM OVER BIT-STRING TRAJECTORIES

The main technical contribution of our work is to derive a correspondence between the collision probability and a discrete sum (which can be interpreted as the partition function of a classical statistical-mechanical model or as the expectation of a Markov chain) and then to derive rigorous upper and lower bounds on the sum. Here, we describe the correspondence along with a brief example for a simple RQC in Fig. 2. We also explain why this correspondence leads us to expect anticoncentration to be achieved after $\Theta(n \log(n))$ gates in most architectures. In Appendix B, we give a more careful derivation of this correspondence, and in the other appendices, we use it to

$$\gamma = \left(\begin{array}{cccccc} \vec{\gamma}^{(0)}, & \vec{\gamma}^{(1)}, & \vec{\gamma}^{(2)}, & \vec{\gamma}^{(3)}, & \vec{\gamma}^{(4)}, & \vec{\gamma}^{(5)} \end{array}\right)$$

$$\text{Weight}(\gamma) = \frac{2}{5} \times \frac{2}{5} \times \frac{2}{5} \times 1 \times \frac{2}{5} = \frac{16}{625}$$

$$\gamma = \left(\begin{array}{cccccc} \vec{\gamma}^{(0)}, & \vec{\gamma}^{(1)}, & \vec{\gamma}^{(2)}, & \vec{\gamma}^{(3)}, & \vec{\gamma}^{(4)}, & \vec{\gamma}^{(5)} \end{array}\right)$$

$$\text{Weight}(\gamma) = \frac{2}{5} \times \frac{2}{5} \times 1 \times 1 \times 1 = \frac{4}{25}$$

$$\text{Collision probability} = Z = \frac{1}{(q+1)^n} \sum_{\gamma} \text{weight}(\gamma)$$

FIG. 2. Two example trajectories for a quantum circuit diagram with $n = 4$ qubits and $s = 5$ gates. Each gate displayed is chosen randomly from the Haar measure over single- or two-qubit unitaries. The collision probability Z is expressed as a weighted sum over trajectories $\gamma = (\vec{\gamma}^{(0)}, \ldots, \vec{\gamma}^{(s)})$, which are length-$(s + 1)$ sequences of assignments ("configurations") of I or S to each of the n qudits. When the input bits to a gate are assigned opposite values, one must be switched at the next configuration in the sequence. These bit flips happen at gates 1, 2, 3, and 5 in the first example and at gates 1 and 2 in the second example. Each bit flip results in a reduction of the weight by a factor 2/5 (when $q = 2$). In the second example, the trajectory reaches one of the fixed-point configurations where all n values agree; this is not the case in the first example. Trajectories that quickly reach a fixed point generally have larger weights and make up most of the contribution to the collision probability.

rigorously prove the upper and lower bounds quoted in Table I.

Recall that we wish to compute the collision probability

$$Z := q^n \mathbb{E}_U \left[p_U(1^n)^2 \right]$$

$$= q^n \mathbb{E}_U \left[\langle 1^n |^{\otimes 2} U^{\otimes 2} | 1^n \rangle \langle 1^n |^{\otimes 2} U^{\dagger \otimes 2} | 1^n \rangle^{\otimes 2} \right], \quad (20)$$

where U is the unitary enacted by the RQC. The Haar measure uniformly covers the unitary group so, intuitively speaking, taking the expectation over application of a Haar-random gate removes much of the bias in the quantum state; we use a technique that allows us to effectively keep track of only n bits of information about the n-qudit state after the application of (two copies of) each Haar-random gate. Instead of 0 or 1, our bits take values I or S, because they are associated with the identity and swap operations on two qudit copies.

In particular, if V is a $q \times q$ Haar-random matrix and σ is an operator on two copies of a q-dimensional Hilbert space, then the quantity $\mathbb{E}_V[V^{\otimes 2} \sigma V^{\dagger \otimes 2}]$ is equal to a linear combination of the identity operation I on the two copies of the Hilbert space and the swap operation S on the two copies of the Hilbert space. Specifically, it is given by

$$\left[\frac{\text{Tr}(\sigma) - q^{-1}\text{Tr}(\sigma S)}{q^2 - 1} \right] I + \left[\frac{\text{Tr}(\sigma S) - q^{-1}\text{Tr}(\sigma)}{q^2 - 1} \right] S. \quad (21)$$

This well-known formula is derived in Appendix B.

By applying the formula to each of the two-qudit Haar-random gates sequentially, the state (which begins in

$|1^n\rangle\langle 1^n|^{\otimes 2}$) evolves as a sum over n-fold tensor products of identity and swap operations. Each of these n-fold tensor products is labeled by an n-bit vector that we call a *configuration* $\vec{v} \in \{I, S\}^n$. For a circuit with s two-qudit gates, each term in the resulting sum is then associated with a length-$(s + 1)$ sequence of configurations $\gamma = (\vec{\gamma}^{(0)}, \ldots, \vec{\gamma}^{(s)})$, which we call a *trajectory*. Each trajectory γ has a certain non-negative coefficient in the sum, allowing us to write

$$Z = \frac{1}{(q+1)^n} \sum_{\gamma} \text{weight}(\gamma) \quad (22)$$

for a fairly simple weight function, described as follows and derived more carefully in Appendix B.

First, the weight for most trajectories is simply 0. In order for a trajectory to have positive weight, it must obey the following rules. If the gate at time step t acts on qudits a and b, then the configuration values $\gamma_a^{(t)}, \gamma_b^{(t)} \in \{I, S\}$ at positions a and b must be equal, either both I or both S. Thus if the values disagree at the previous time step, i.e., $\gamma_a^{(t-1)} \neq \gamma_b^{(t-1)}$, one of the bits must be flipped during the transition from $\vec{\gamma}^{(t-1)}$ to $\vec{\gamma}^{(t)}$. If the values at positions a and b already agree at time step $t - 1$, they must remain unchanged from time step $t - 1$ to time step t. Moreover, the bit values at the other $n - 2$ positions must also remain unchanged from time step $t - 1$ to time step t.

For trajectories that obey these rules, the weight begins at 1 but each time a bit flip occurs, the weight is reduced by a factor of 2/5 for qubits, or $q/(q^2 + 1)$ for general local dimensions. Thus, the most significant terms in the weighted sum are the terms with the fewest bit flips along the trajectory.

The expression for Z as a weighted sum can alternatively be interpreted as a partition function for an Ising-like classical statistical-mechanical model, since it is a weighted sum over "spin" configurations for spins with two possible values, or it can be interpreted as the expectation of a certain quantity over a simple Markov chain that generates the sequence $(\vec{\gamma}^{(0)}, \ldots, \vec{\gamma}^{(s)})$. We take the latter approach in our application of the method to prove our upper and lower bounds. For an example of two trajectories for a simple RQC, along with a calculation of their weight, see Fig. 2.

The correspondence given in Eq. (22) is powerful because we have a good sense of what to expect from the weighted sum over trajectories and we can draw conclusions that are not obvious from the definition of the collision probability itself. For example, we can straightforwardly analyze the infinite circuit-size limit. In this limit, each positive-weight trajectory γ will be forced to keep flipping bits (each time a two-qudit gate acts on a disagreeing pair of bits) until it reaches a fixed point, either I^n or S^n, in which case bits can no longer be flipped since all the bits agree. Let $Q(x)$ be the total weight of all trajectories that begin at a configuration with x S assignments and $n-x$ I assignments. At some point in the circuit, a disagreeing pair of bits will be acted upon by a gate and one of the bits must flip, sending the number of S assignments either to $x-1$ or $x+1$ and reducing the weight by $q/(q^2+1)$. Since there are an infinite number of gates, the following recursion relation must be obeyed:

$$Q(x) = \frac{q}{q^2+1}\left[Q(x-1) + Q(x+1)\right], \quad (23)$$

which, by imposing the boundary conditions $Q(0) = Q(n) = 1$, has the unique solution

$$Q(x) = \frac{q^x + q^{-x}}{q^n + 1}. \quad (24)$$

Moreover, for each x, there are $\binom{n}{x}$ configurations each contributing weight $Q(x)$, so

$$\lim_{s \to \infty} Z = \frac{\sum_{x=0}^{n}\binom{n}{x}(q^x + q^{-x})}{(q+1)^n(q^n+1)} = \frac{2}{q^n+1} = Z_H, \quad (25)$$

reproducing the value Z_H that would be obtained if the RQC was one large $q^n \times q^n$ Haar-random transformation instead of a series of $q^2 \times q^2$ two-qudit gates. [The fact that a $q^n \times q^n$ Haar-random transformation yields Z_H is a direct consequence of Eq. (21) with the substitution $q \to q^n$.] This conclusion makes sense since a random circuit with an infinite number of 2-local Haar-random gates should enact a global Haar-random transformation.

When the circuit size is a finite number s, we have $Z > Z_H$, corresponding to the fact that many trajectories have not yet reached a fixed point and are overweighted compared to their contribution to Z_H. As the circuit size increases, more of the trajectories get closer to the fixed point and Z approaches Z_H. The point at which anticoncentration is achieved is intimately connected with the point at which most of the weight can be accounted for by trajectories that have reached a fixed point. A depiction of this process at $n = 60$ is given in Fig. 3.

Our quantitative challenge is to understand, for a certain RQC architecture, how quickly these trajectories approach the fixed points, and consequently how quickly Z approaches Z_H, as the circuit size increases. Recall that we define the *anticoncentration size* s_{AC} to be the circuit size (as a function of the number of qudits n) needed for Z to be only a constant factor larger than Z_H. Perhaps surprisingly, we find in multiple architectures that $s_{AC} = \Theta(n \log(n))$, corresponding to only $\Theta(\log(n))$ gates per qudit. We can explain this observation heuristically by generating trajectories γ at random with probability proportional to weight(γ) (in the statistical-mechanical interpretation, this corresponds to drawing samples from the thermal distribution). For typical trajectories generated in this fashion, each additional layer of $\Theta(n)$ gates will cause the trajectory to move a constant fraction of the way closer to terminating at a fixed point. Since trajectories typically begin on the order of n bit flips away from the fixed point [i.e., the initial configuration typically has $\Theta(n)$ I assignments and $\Theta(n)$ S assignments], $\Omega(\log(n))$ layers are necessary and sufficient for typical trajectories to get within a constant distance from the fixed point.

This heuristic statement is perhaps confirmed most clearly in the complete-graph architecture, where qudit pairs are chosen uniformly at random. Here, let $x \ll n$ and suppose that the current configuration at time step t has value S at x of the n positions and value I at the other $n-x$ positions. If we perform gates on $n/2$ random pairs of qudits, we expect roughly x of those pairs to couple an I value with an S value. Each time this happens, a bit must be flipped and there is an opportunity for the trajectory to move closer to the fixed point I^n. Thus, we expect the number of S values in the configuration at time step $t + n/2$ to have decreased by an amount proportional to x. After $\Theta(n \log(n))$ gates, we expect the trajectory to be at (or very close to) the fixed point I^n with high probability. Fewer gates would leave most trajectories too far from the fixed point for anticoncentration to have been reached. In Fig. 3, we illustrate the convergence of typical trajectories and the correspondent convergence of Z for the complete-graph architecture at $n = 60$.

We prove that a similar situation occurs even if the gates are arranged in a 1D fashion and we fully expect that this situation applies for nearly all natural [46] architectures, including circuits on D-dimensional lattices for $D > 1$. We formalize this in Conjecture 1. We believe Conjecture 1 first because anticoncentration should intuitively only be faster when the circuit becomes more connected and the
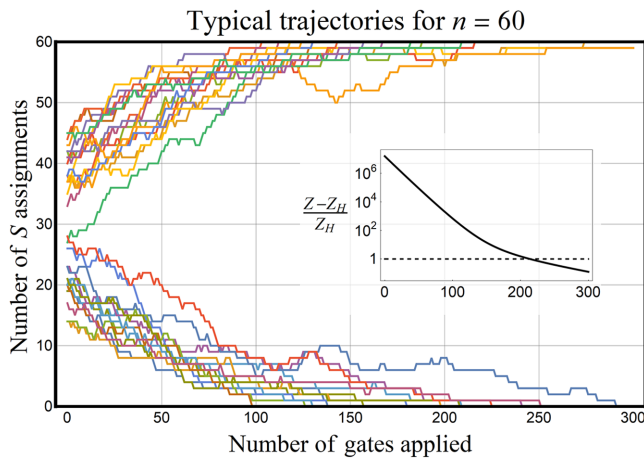
FIG. 3. Thirty trajectories generated randomly for the complete-graph architecture at $n = 60$. A trajectory γ is chosen with probability proportional to weight(γ) in the $s \to \infty$ limit and then the number of S assignments (out of 60) are plotted for the first 300 time steps. The trajectories rapidly approach either the fixed point I^n with 0 S assignments, or the fixed point S^n with 60 S assignments but not all have reached the fixed point within 300 time steps. The distance of a typical trajectory from the nearest fixed point decays exponentially with time, with characteristic time scale $\Theta(n)$. Thus, it takes $\Theta(n \log(n))$ gates for most typical trajectories to have reached the fixed point. In the inset, it is shown that as trajectories approach the fixed points, the collision probability Z (which can be efficiently numerically calculated for the complete-graph architecture) approaches Z_H. Anticoncentration is defined as the point where it falls beneath $2Z_H$ (dashed line), which occurs at $s = 214$ for $n = 60$.

1D architecture is perhaps the least connected a natural architecture can be, as it takes $\Omega(n^2)$ gates for information to travel across the diameter of the qudit array. Second, the above intuitive argument about the convergence of typical trajectories to a fixed point in $O(n \log(n))$ gates should apply to any natural architecture. Specifically, if we choose a configuration with x S assignments at random and we apply a layer of $\Theta(n)$ two-qudit gates, with high probability we will have formed $\Theta(x)$ disagreeing pairs and moved the trajectory a constant fraction of the way to the nearest fixed point. The difficulty in proving Conjecture 1 lies in characterizing what happens in the low-probability event that this is not the case.

Indeed, our rigorous proofs for the 1D and complete-graph architectures have to deal with the fact that it is not sufficient to examine only typical trajectory behavior. In particular, the collective contribution of trajectories at the tails of what is allowed are tricky to bound. Nonetheless, heuristic reasoning about typical trajectory behavior ultimately gives accurate predictions about the collision probability in these cases.

The rigorous bounds are provided in the appendices. In the 1D case, the proof associates each trajectory with

a configuration of domain walls on a 2D lattice [of size $n \times (d + 1)$, where d is the circuit depth] and bounds their total contribution combinatorially, similar to the method employed in Refs. [2,3,17,32]. The main intuition is that trajectories that have not reached a fixed point must have domain walls that penetrate through the entire depth of the circuit and thus receive weight that decreases exponentially with the depth as $[q/(q^2 + 1)]^d$. If one accounts for the total number of possible domain walls of this type, which scales like $n2^d$, one finds that $d = O(\log(n))$ is sufficient for the overall contribution to be small. We use a similar domain-wall counting method to produce a tight lower bound.

In the complete-graph case, we present a much different and novel approach. For each trajectory γ, we define $R[\gamma]$ as the "reduced" trajectory that results from removing consecutive duplicates of the same configuration. Note that the weight of γ depends only on the length of $R[\gamma]$ (the number of bit flips). Long subsequences of consecutive duplicates occur when the randomly chosen gates repeatedly act on pairs of qudits that are already assigned the same value by the configuration, an outcome that becomes more likely as the trajectory approaches a fixed point. For each ψ, we can condition on $R[\gamma] = \psi$ and examine the probability distribution over the length of γ (i.e., the number of consecutive duplicates plus the number of bit flips). We use a Chernoff bound to upper bound the probability that the length of γ is greater than a certain quantity. We then use a generalization of the recursive method that yields Eq. (24) to perform the weighted sum over all reduced trajectories ψ. In the appendix, we provide a more detailed proof summary prior to the full proof.

VI. OUTLOOK

In a quantum computer, quantum information is ultimately accessed by making measurements of the output state and obtaining samples from the associated output distribution over measurement outcomes. In many applications, it is desirable to choose our quantum computation completely at random, the only constraint being the arrangement of the different gates, and thus it is important to characterize the output distribution over measurement outcomes in RQCs and how it depends on the underlying circuit architecture.

One feature of the output distribution is that, for very shallow circuits, there are a relatively small number of very "heavy" measurement outcomes that are exponentially more likely than average to be obtained, a fact that inhibits the design of certain classical simulation algorithms but also, in other cases, prevents potential proofs that no good simulation algorithms exist. As the circuit gets deeper, the probability mass gradually anticoncentrates and eventually becomes fairly well spread out over all possible measurement outcomes. We develop a framework to quantitatively

understand this situation; we map the anticoncentration process to the equilibration of a simple stochastic process (an alternative interpretation of the stochastic process is the partition function of a statistical-mechanical model). The stochastic process allows for effective qualitative reasoning but also produces sharp quantitative anticoncentration upper and lower bounds.

Both sides of our bounds have meaningful and surprising takeaways. On the one hand, the fact that only $O(n \log(n))$ gates are needed to achieve anticoncentration in geometrically local and nonlocal architectures contradicts the intuition that anticoncentration should not occur until information has had time to spread across the entire system. In fact, up to a constant factor, the anticoncentration time does not appear to be sensitive to the exact connectivity structure of the circuit. While we only rigorously consider two architectures, our work provides strong evidence that any natural architecture anticoncentrates in $O(n \log(n))$ gates [which typically corresponds to $O(\log(n))$ depth]. In cases where anticoncentration is a desirable property, our work gives explicit bounds on how many gates are needed and the fact that this number is relatively small will come as welcome news in practical situations where the gates are noisy or otherwise costly to implement.

On the other hand, by showing that $\Omega(n \log(n))$ gates are necessary for anticoncentration (and computing the optimal constant prefactor in our two specific scenarios), we clear up some confusion about very shallow circuits. An increase in the depth causes the anticoncentration process to begin but our lower bound implies that the phenomenon of very heavy measurement outcomes will remain for any architecture of constant depth. Even the 2D circuits of depth $O(\sqrt{\log(n)})$ [for which the light-cone volume is $O(\log(n))$] considered in Ref. [32] cannot be anticoncentrated, as had been speculated in that work.

We conclude with some other specific open problems inspired by our work:

(a) We prove that the anticoncentration size is $\Theta(n \log(n))$ for the 1D and complete-graph architectures. We believe this is true for most other natural architectures and formally conjecture in Conjecture 1 that this follows from our "regularly connected" definition.

(b) A sharp anticoncentration analysis for 2D and higher-dimensional geometrically local architectures would be particularly valuable since, unlike in 1D, $\Theta(\log(n))$-depth 2D circuits can perform universal quantum computation [indeed, $\Omega(1)$-depth is sufficient [47]] and 2D circuits form the basis for Google's 2019 quantum computational supremacy experiment [6].

(c) We suspect that the constant prefactor of $[2 \log(q^2 + 1)]^{-1}$ in the general lower bound in Theorem 4 could

be improved. What is its optimal value? That is, can we show an improved general lower bound and then find a RQC architecture $\mathcal{A}_{\text{fast}}$ that has a matching upper bound? This would show that $\mathcal{A}_{\text{fast}}$ is the fastest anticoncentrator. A candidate for $\mathcal{A}_{\text{fast}}$ is the architecture where each layer of $n/2$ gates is formed by choosing a random partition of the n qudits into $n/2$ pairs.

(d) Are there other problems involving second-moment calculations over RQCs where our techniques would produce sharp upper and lower bounds? One such problem could be the 2-design time for RQCs in various architectures.

APPENDIX A: DEFINITIONS: RANDOM QUANTUM CIRCUIT ARCHITECTURES, ANTICONCENTRATION, AND THE COLLISION PROBABILITY

1. Random quantum circuits (RQCs)

Here, we establish some precise definitions for the terms in this paper. Throughout, we consider systems of n qudits of local Hilbert space dimension q, with basis states $\{|1\rangle, |2\rangle, \ldots, |q\rangle\}$. Loosely speaking, a quantum circuit is a sequence of unitary transformations called gates, which each typically involve only a few of the n qudits, acting on the initial state $|1\rangle^{\otimes n} \equiv |1^n\rangle$. Formally, we let a *quantum circuit diagram* of circuit size s be specified by a length-s sequence $(A^{(1)}, \ldots, A^{(s)})$ of nonempty subsets of $[n] := \{1, 2, \ldots, n\}$, indicating for each gate which qudits participate in that gate. Since we consider circuits consisting only of two-qudit gates, we require $|A^{(t)}| = 2$ for all t. We also make the assumption that the circuit begins with a single-qudit gate on each of the n qudits at the beginning of the circuit, without counting these n gates toward

the circuit size. This sequence can be turned into a diagram as in Fig. 2 (ignoring the overlaid I and S), where the gate sequence is $(\{1, 2\}, \{2, 3\}, \{1, 2\}, \{3, 4\}, \{2, 3\})$. Note that the single-qudit unitaries are each displayed with the symbol U but will not necessarily be the same unitary. The circuit *depth d* of a circuit diagram is the minimum number of layers of nonoverlapping gates needed to implement all s gates in the circuit or, formally, the smallest integer such that there exists a sequence $0 = s_0 < s_1 < s_2 < \ldots < s_d = s$ where $A^{(t)} \cap A^{(t')} = \emptyset$ whenever $s_j < t < t' \leq s_{j+1}$.

Once a circuit diagram is chosen, a *quantum circuit instance* is generated by additionally specifying a length-$s + n$ sequence of unitary matrices $(U^{(-n)}, \ldots, U^{(-1)}, U^{(1)}, \ldots, U^{(s)})$, where $U^{(-j)}$ is a $q \times q$ (single-qudit) matrix for each $j = 1, \ldots, n$ and $U^{(t)}$ is a $q^2 \times q^2$ (two-qudit) matrix for each $t = 1, \ldots, s$. We denote the global $q^n \times q^n$ unitary operator implemented by the circuit by U, where

$$U := U_{A^{(s)}}^{(s)} U_{A^{(s-1)}}^{(s-1)} \ldots U_{A^{(2)}}^{(2)} U_{A^{(1)}}^{(1)} U_{\{1\}}^{(-1)} \ldots U_{\{n\}}^{(-n)}, \quad \text{(A1)}$$

with V_X indicating the action of the $q^{|X|} \times q^{|X|}$ unitary V on the qudits in subregion $X \subset [n]$ tensored with the identity operation on the qudits in the complement of X.

In this work, we always assume that projective computational basis measurements are performed on all n qudits at the end of the circuit. Thus, a quantum circuit instance U has a corresponding classical probability distribution p_U over possible measurement outcomes $x \in [q]^n$, as follows:

$$p_U(x) = \left| \langle x | U | 1^n \rangle \right|^2. \quad \text{(A2)}$$

RQCs refer to situations when, once a circuit diagram has been fixed, the actual unitary gates $U^{(t)}$ that determine the circuit instance are each randomly chosen independently from some distribution over the unitary group. In this paper, we always take this distribution to be the Haar measure, but since our techniques rely on calculating expectations over quantities with only two copies of each $U^{(t)}$, our results also apply when the gates are drawn from any 2-design, such as the Clifford group. Note that Google's quantum computational supremacy experiment [6] has drawn gates from another distribution that is not a 2-design. Heuristically speaking, as long as the distribution lacks any bias or symmetries, we expect properties such as anticoncentration to be the same as in the Haar-random case.

2. Random quantum circuit architectures

An *architecture* for RQCs is simply a procedure for choosing a circuit diagram. Formally, we define it to be a (possibly randomized) classical algorithm that, given parameters n and s, computes a circuit diagram of size s on n qudits. Given an architecture and parameters n and s, we let the expectation of some quantity Q, denoted $\mathbb{E}_U[Q]$,

refer to the expectation over the process of first choosing a circuit diagram according to the architecture and then choosing a circuit instance by randomly generating each gate in the circuit diagram independently from the Haar measure. Next, we define the two architectures that we consider.

Definition 1 (Complete-graph architecture). *Circuit diagrams of size s on n qudits are generated by choosing s gates each uniformly at random from the set of all two-qudit gates, i.e., $A^{(t)}$ is chosen uniformly from $\{\{a, b\} : a, b \in [n], a \neq b\}$.*

Note that if it could be guaranteed that every qudit would eventually participate in at least one gate, the distribution over circuit instances would be equivalent if we omitted the first layer of n single-qudit gates (defined to be part of every architecture), a fact that follows from the invariance of the Haar measure; the single-qudit gates could be absorbed into the two-qudit Haar-random gates that act directly before or after without changing the distribution over the two-qudit gates. However, in the complete-graph architecture there is a chance that a qudit does not participate in any two-qudit gates, although for sufficiently large circuit size the probability of this vanishes.

Definition 2 (1D architecture). *Assume that n is even and that $d := 2s/n$ is an integer. The circuit diagram of size s on n qudits is generated by alternating between the two types of layers of $n/2$ nonoverlapping nearest-neighbor two-qudit gates on a ring. That is, for each $t = 1, \ldots, n/2$, if j is even, then $A^{(t+jn/2)} = \{2t-1, 2t\}$, and if j is odd, then $A^{(t+jn/2)} = \{2t, 2t+1\}$, where index n is identified with index 0 to enforce periodic boundary conditions.*

3. Collision probability and anticoncentration

Anticoncentration is a concept that describes a classical probability distribution for which the probability mass is not too concentrated onto a small number of outcomes of the random variable. The uniform distribution is the ultimate anticoncentrated distribution, as the probability mass is allocated evenly over every possible outcome, but we would still like the term "anticoncentrated" to apply to some nonuniform distributions if the probability mass is fairly well spread over many of the outcomes. There are multiple ways to make this quantitative. For the purposes of this paper, we choose one way—the collision probability—that mirrors previous work on anticoncentration of quantum circuit outputs and suffices for the applications we discuss in Sec. I.

Let X be a discrete random variable and let M be the set of outcomes of X. We can form another random variable p, where p is equal to $\Pr[X = x]$ for an x chosen uniformly at random from M. Since $\sum_x \Pr[X = x] = 1$, we

have $\mathbb{E}[p] = 1/|M|$ no matter how X is distributed. We define the collision probability for X to be

$$
\begin{aligned}
Z &:= \sum_{x \in M} \Pr[X = x]^2 \\
&= |M|\mathbb{E}[p^2] \\
&= |M|\mathrm{var}(p) + |M|^{-1},
\end{aligned}
\tag{A3}
$$

which is the probability that two identical independent copies of X will be equal to each other—hence *collision* probability. If the distribution over X is the uniform distribution, then the distribution over p is the point distribution on the value $|M|^{-1}$, the collision probability takes its minimal value $Z = |M|^{-1}$, and $\mathrm{var}(p) = 0$. If X is nonuniform but still somewhat anticoncentrated, then p will not always be $|M|^{-1}$ but it will usually be close and this will be reflected by a collision probability that is greater, but not too much greater, than $|M|^{-1}$. Formally, we make the following definition.

Definition 3 (Anticoncentrated). *We say that a random variable X over a set M of outcomes is α-anticoncentrated for $0 < \alpha \le 1$ if*

$$
Z := \sum_{x \in M} \Pr[X = x]^2 \le \frac{1}{|M|\alpha}.
\tag{A4}
$$

Thus a distribution is 1-anticoncentrated if and only if it is the uniform distribution.

In our setting, the random variable X is the measurement outcome of a RQC instance, which is distributed according to the distribution p_U over the outcome set $[q]^n$. Example distributions of p_U for RQC outputs in the uniform, the nonuniform but still anticoncentrated and the not anticoncentrated case are shown in the sketch in Fig. 1. A RQC architecture for specified n and s is understood as an ensemble over many different U, only some of which will have output distributions p_U that are α-anticoncentrated for a certain choice of α. We would like to say that the architecture as a whole is anticoncentrated if typical circuit instances drawn from the architecture are anticoncentrated, acknowledging that not every instance will be. We also require this to hold for the same constant α as n increases, with s increasing like some function $s(n)$. Formally, we accomplish this by averaging the collision probability over the random circuit instance, as follows.

Definition 4 (Anticoncentrated RQC architecture). *We say that a RQC architecture is α-anticoncentrated for $0 < \alpha \le 1$ at circuit size $s(n)$ if there exists n_0 such that whenever*

$n \ge n_0$

$$
Z := \mathbb{E}_{U}\left[\sum_{x \in [q]^n} p_U(x)^2 \right] \le (\alpha q^n)^{-1},
\tag{A5}
$$

where $\mathbb{E}_U$ denotes drawing circuit instances according to the architecture over n qudits with circuit size $s(n)$. Generally, we say that the architecture is anticoncentrated at size $s(n)$ if there exists a constant $\alpha > 0$ independent of n for which it is α-anticoncentrated at that size.

RQC architectures for which every qudit experiences at least one gate, which includes all the architectures introduced above, will have a symmetry over the q^n measurement outcomes in the sense that the quantity $p_U(x)$ is distributed identically (over circuit instances) for every x. In this case, each term in the sum in Eq. (A5) will have the same contribution and we can simply write

$$
Z = q^n \mathbb{E}_{U}\left[p_U(1^n)^2 \right].
\tag{A6}
$$

The anticoncentration of an architecture implies that most of the instances drawn from that architecture have good anticoncentration properties. Given an architecture at a certain size and a bound on its collision probability $Z \le \alpha^{-1} q^{-n}$, we can use Markov's inequality to assert that at least a $1 - \beta$ fraction of instances have collision probability at most $q^{-n}\left[1 + (\alpha^{-1} - 1)\beta^{-1}\right]$. In practice, we expect the collision probability of individual instances to be even more clustered near the mean collision probability than this analysis indicates, but proving that this is the case would seem to require computing higher moments such as $\mathbb{E}_U[p_U(1^n)^k]$ for $k > 2$.

As discussed in the main text, an important implication of an α-anticoncentrated architecture is that for any β with $0 \le \beta \le 1$ and sufficiently large n,

$$
\Pr_{U}[p_U(x) \ge \beta q^{-n}] \ge (1 - \beta)^2 \alpha,
\tag{A7}
$$

which follows directly from the Paley-Zygmund inequality. This inequality indicates that whenever an architecture is anticoncentrated, at least a constant fraction of the outcomes will be allocated an amount of mass that is within a constant factor β of the mean mass; it cannot be the case that all but a vanishing fraction of the outcomes are allocated a vanishing fraction of the mean mass.

APPENDIX B: FRAMEWORK FOR ANALYSIS: RANDOM QUANTUM CIRCUITS AS A STOCHASTIC PROCESS

This appendix gives more details on the correspondence discussed in Sec. V of the main text. The key idea in our analysis of the collision probability of RQCs is to perform
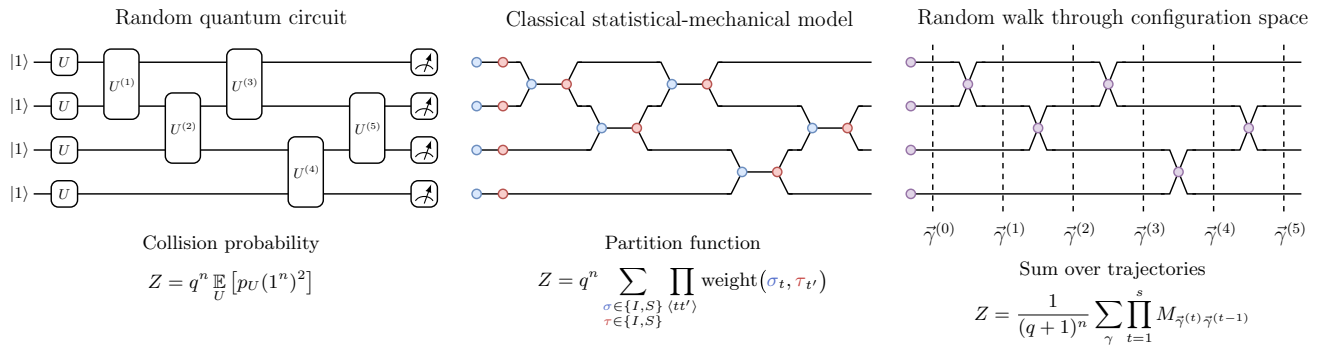
FIG. 4.   A diagram depicting the equivalent ways to interpret the expected value of the collision probability for RQCs. Left: a RQC of size five. Middle: the reinterpretation as the partition function of a classical statistical-mechanical model with local Ising-like spins. Right: another interpretation as a stochastic process of evolving configurations.

the Haar expectation over each local unitary individually. This is possible due to explicit formulas for expectations under action by a Haar-random unitary. We use these formulas to reexpress the collision probability, originally an integral over many continuously varying unitary matrices drawn from the Haar measure, as a weighted discrete sum, which is then analyzed using combinatorial and stochastic methods. This weighted sum can also be interpreted as the partition function of a classical statistical-mechanical Ising-like model or as the expectation value of a simple stochastic process. Figure 4 depicts these equivalent representations of the problem. In this appendix, we explain this method and derive the important formulas that will apply generally for any RQC architecture, which are then used in Appendices C, D, and E to prove our main results.

1. Averaging individual unitaries over the Haar measure

The quantity of interest for anticoncentration is the expected collision probability, which is proportional to a second moment over choice of unitary operator U, as illustrated in the following equation, where we recall that $|1^n\rangle\langle 1^n|^{\otimes 2}$ is two copies of the circuit input state:

$$Z := q^n \mathop{\mathbb{E}}_{U}\left[\left(\langle 1^n|\, U\, |1^n\rangle \langle 1^n|\, U^\dagger\, |1^n\rangle\right)^2\right]$$

$$= q^n \mathrm{Tr}\left[\mathop{\mathbb{E}}_{U}\left[U^{\otimes 2}|1^n\rangle\langle 1^n|^{\otimes 2}U^{\dagger\otimes 2}\right]|1^n\rangle\langle 1^n|^{\otimes 2}\right]. \quad \text{(B1)}$$

Moreover, for a fixed quantum circuit diagram, the unitary U is given by Eq. (A1) as a product of single-qudit unitaries $U^{(-j)}$ acting on qudit j for $j = 1,\ldots,n$ and two-qudit unitaries $U^{(t)}$ acting on some pair of qudits $A^{(t)} \subset [n]$ for $t = 1,\ldots,s$. Each unitary is independently chosen according to the Haar measure and its expectation can be evaluated

separately. Let

$$M^{(t)}(\rho) := \mathop{\mathbb{E}}_{U^{(t)}}\left[U^{(t)}_{A^{(t)}}{}^{\otimes 2}\rho\, U^{(t)}_{A^{(t)}}{}^{\dagger\otimes 2}\right]. \quad \text{(B2)}$$

Then, we can write

$$\mathop{\mathbb{E}}_{U}\left[U^{\otimes 2}|1^n\rangle\langle 1^n|^{\otimes 2}U^{\dagger\otimes 2}\right]$$

$$= M^{(s)} \circ M^{(s-1)} \circ \cdots \circ M^{(1)}$$

$$\circ M^{(-1)} \circ \ldots \circ M^{(-n)}\left(|1^n\rangle\langle 1^n|^{\otimes 2}\right). \quad \text{(B3)}$$

When an architecture is itself a mixture over randomly chosen circuit diagrams, such as the complete-graph architecture, the overall quantity $\mathbb{E}_U[U^{\otimes 2}|1^n\rangle\langle 1^n|^{\otimes 2}U^{\dagger\otimes 2}]$ is a mixture over terms of the above form.

The remainder of this subsection illustrates how the action of $M^{(t)}$ can be evaluated, ultimately allowing us to arrive at the expression for Z given in Eq. (B18). In the other subsections of this section, we explain how that equation can be interpreted as a partition function of a classical statistical-mechanical model or as the expectation over simple stochastic process.

When the local unitaries are drawn from the Haar measure (or any exact 2-design), the expression $M^{(t)}(\rho)$ can be evaluated in a simple way. Generally, for σ a $q^2 \times q^2$ Hermitian operator, and with V chosen from the Haar measure over the set of $q \times q$ unitaries, we define

$$M(\sigma) := \mathop{\mathbb{E}}_{V}\left[V^{\otimes 2}\sigma V^{\dagger\otimes 2}\right] \quad \text{(B4)}$$

and observe that, for any unitary W and any σ,

$$M(\sigma)W^{\otimes 2} = \mathop{\mathbb{E}}_{V}\left[V^{\otimes 2}\sigma\left(W^\dagger V\right)^{\dagger\otimes 2}\right] = \mathop{\mathbb{E}}_{V}\left[(WV)^{\otimes 2}\sigma V^{\dagger\otimes 2}\right]$$

$$= W^{\otimes 2}M(\sigma), \quad \text{(B5)}$$

where the second equality follows from the invariance of the Haar measure under the substitution $V \to WV$. A mathematical fact from Schur-Weyl duality (see Ref. [48]) is that any operator on k copies of a system that commutes with $W^{\otimes k}$ for any unitary W must be a linear combination of permutation operators over the k systems. Here, we have $k = 2$ and thus the only permutation operators are the identity operation I and the swap operation S, which can be defined as the operator satisfying $S |\psi\rangle \otimes |\phi\rangle = |\phi\rangle \otimes |\psi\rangle$ for any $|\psi\rangle, |\phi\rangle$. Letting $M(\sigma) = \alpha I + \beta S$, we make the following calculations:

$$\mathrm{Tr}[M(\sigma)] = \mathrm{Tr}(\sigma) = \alpha q^2 + \beta q, \tag{B6}$$

$$\mathrm{Tr}[M(\sigma)S] = \mathrm{Tr}(\sigma S) = \alpha q + \beta q^2, \tag{B7}$$

which determine α and β and allow us to write

$$M(\sigma) = \frac{\mathrm{Tr}(\sigma) - q^{-1}\mathrm{Tr}(\sigma S)}{q^2 - 1} I + \frac{\mathrm{Tr}(\sigma S) - q^{-1}\mathrm{Tr}(\sigma)}{q^2 - 1} S. \tag{B8}$$

The unitaries $U^{(-j)}$ are $q \times q$ (single-qudit) that act on qudit j. Two copies of the input state on qudit j is $|1\rangle\langle 1|_{\{j\}}^{\otimes 2}$. Denote two copies of the input state on the other $n - 1$

qudits by $\rho_{[n]\backslash\{j\}}$. Using Eq. (B8), we then find that

$$M^{(-j)}\left(\rho_{[n]\backslash j} \otimes |1\rangle\langle 1|_{\{j\}}^{\otimes 2}\right)$$
$$= \rho_{[n]\backslash\{j\}} \otimes \frac{1}{q(q + 1)} (I + S)_{\{j\}}, \tag{B9}$$

meaning that $M^{(-j)}$ simply replaces the state on qudit j as a uniform sum over operators I and S. Hence

$$M^{(-1)} \circ \cdots \circ M^{(-n)}\left(|1^n\rangle\langle 1^n|^{\otimes 2}\right)$$
$$= \bigotimes_{j=1}^{n}\left(\frac{1}{q(q + 1)} (I + S)_{\{j\}}\right)$$
$$= \frac{1}{q^n(q + 1)^n} \sum_{\vec{\gamma} \in \{I,S\}^n} \bigotimes_{j=1}^{n} \gamma_j. \tag{B10}$$

We call each $\vec{\gamma} \in \{I, S\}^n$ a *configuration*. The above equation states that the expected value of two copies of the state after application of all the single-qudit unitaries is precisely a uniform sum over all identity and/or swap configurations of the n sites.

Now, we need to examine the action of $M^{(t)}$ for $t > 0$. In this case, the unitaries are $q^2 \times q^2$ and they act on the qudit pair $A^{(t)}$. We can use Eq. (B8) by replacing $q \to q^2$ and sending $I \to I \otimes I$, the identity operation on two copies of two qudits, and $S \to S \otimes S$, the swap operation on two copies of two qudits. We assume that the input state is a product state $\rho_{[n]\backslash A^{(t)}} \otimes \rho_{A^{(t)}}$ and see that

$$M^{(t)}(\rho_{[n]\backslash A^{(t)}} \otimes \rho_{A^{(t)}})$$
$$= \rho_{[n]\backslash A^{(t)}} \otimes \left[\frac{\mathrm{Tr}(\rho_{A^{(t)}}) - q^{-2}\mathrm{Tr}\left(\rho_{A^{(t)}}(S \otimes S)\right)}{q^4 - 1}(I \otimes I)_{A^{(t)}} + \frac{\mathrm{Tr}\left(\rho_{A^{(t)}}(S \otimes S)\right) - q^{-2}\mathrm{Tr}(\rho_{A^{(t)}})}{q^4 - 1}(S \otimes S)_{A^{(t)}}\right]. \tag{B11}$$

Since the two qudit gates act after the single-qudit gates, the input state to $M^{(t)}$ will always be a sum of tensor products of I and S, so we only need to evaluate the above expression when $\rho_{A^{(t)}}$ is either $I \otimes I, I \otimes S, S \otimes I$, or $S \otimes S$. Doing so, we arrive at

$$M^{(t)}\left(\rho_{[n]\backslash A^{(t)}} \otimes (I \otimes I)_{A^{(t)}}\right) = \rho_{[n]\backslash A^{(t)}} \otimes (I \otimes I)_{A^{(t)}}, \tag{B12}$$

$$M^{(t)}\left(\rho_{[n]\backslash A^{(t)}} \otimes (S \otimes S)_{A^{(t)}}\right) = \rho_{[n]\backslash A^{(t)}} \otimes (S \otimes S)_{A^{(t)}}, \tag{B13}$$

$$M^{(t)}\left(\rho_{[n]\backslash A^{(t)}} \otimes (I \otimes S)_{A^{(t)}}\right) = M^{(t)}\left[\rho_{[n]\backslash A^{(t)}} \otimes (S \otimes I)_{A^{(t)}}\right]$$
$$= \rho_{[n]\backslash A^{(t)}} \otimes \left[\frac{q}{q^2 + 1}(I \otimes I)_{A^{(t)}} + \frac{q}{q^2 + 1}(S \otimes S)_{A^{(t)}}\right]. \tag{B14}$$

Thus, if ρ is a linear combination of configurations in $\{I, S\}^n$, $M^{(t)}(\rho)$ will also be a linear combination of configurations, with coefficients that transform linearly under application of $M^{(t)}$. For configurations $\vec{\gamma}, \vec{\nu} \in \{I, S\}^n$, we let $M^{(t)}_{\vec{\nu}\vec{\gamma}}$ be the matrix element of this linear transformation, defined such that

$$M^{(t)} \left(\bigotimes_{j=1}^{n} \gamma_j \right) = \sum_{\vec{v} \in \{I,S\}^n} M^{(t)}_{\vec{v}\vec{\gamma}} \bigotimes_{j=1}^{n} v_j. \tag{B15}$$

Suppose that $U^{(t)}$ acts on qudits $A^{(t)} = \{a,b\} \subset [n]$. Then, from Eqs. (B12), (B13), (B14), we have

$$M^{(t)}_{\vec{v}\vec{\gamma}} = \begin{cases} 1, & \text{if } \gamma_a = \gamma_b \text{ and } \vec{\gamma} = \vec{v}, \\ \dfrac{q}{q^2+1}, & \text{if } \gamma_a \neq \gamma_b \text{ and } v_a = v_b \text{ and } \gamma_c = v_c \ \forall c \in [n] \setminus \{a,b\}, \\ 0, & \text{otherwise.} \end{cases} \tag{B16}$$

Importantly, $M^{(t)}_{\vec{v}\vec{\gamma}}$ is always non-negative. The way to think about the above equation is to note three things. First, the input configuration $\vec{\gamma}$ and the output configuration $\vec{v}$ must agree on all indices that are not involved in the gate, i.e., for all indices $c \notin \{a,b\}$; otherwise, the matrix element is 0. Second, if the two input values involved in the gate agree, i.e., if $\gamma_a = \gamma_b$ then $v_a = v_b = \gamma_a = \gamma_b$ must hold (in which case the matrix element is 1); otherwise, it is 0. Third, if the two input values disagree, then one of them must be flipped so that the two output values agree [in which case the matrix element is reduced to $q/(q^2+1)$]; otherwise, it is 0.

Note also that

$$\text{Tr}\left[\left(\bigotimes_{j=1}^{n} \gamma_j \right) |1^n\rangle\langle 1^n|^{\otimes 2} \right] = 1 \tag{B17}$$

for all $\vec{\gamma} \in \{I,S\}^n$. Thus, from Eq. (B1), we find

$$Z = \frac{1}{(q+1)^n} \sum_{\gamma \in \{I,S\}^{n \times (s+1)}} \prod_{t=1}^{s} M^{(t)}_{\vec{\gamma}^{(t)} \vec{\gamma}^{(t-1)}},$$

$$=: \frac{1}{(q+1)^n} \sum_{\gamma} \text{weight}(\gamma), \tag{B18}$$

which is the expression quoted in Eq. (22) from the main text. In the above equation, the sum is over length-$(s+1)$ sequences of configurations, which we call a *trajectory* $\gamma = (\vec{\gamma}^{(0)}, \ldots, \vec{\gamma}^{(s)})$, and the weight of each term is given by the product of the matrix elements for each step in the trajectory. This final equation is depicted graphically in the right-hand part of Fig. 4.

2. Collision probability as statistical-mechanical partition function

The expression for the collision probability in Eq. (B18) can be interpreted as a partition function for a classical statistical-mechanical model by thinking of each $\gamma_j^{(t)}$ as an Ising spin variable with the association $\{I, S\} \leftrightarrow \{+1, -1\}$. A trajectory γ is then a configuration of the Ising spins

and Z is a weighted sum over all the spin configurations. Moreover, the weight is always non-negative and is given by a product of factors $M^{(t)}_{\vec{\gamma}^{(t)} \vec{\gamma}^{(t-1)}}$ that can be determined by examining a small number of the spin values. This means that the energy functional over spin configurations of the classical Ising-like model is always real and can be broken up into local terms that depend on the local dimension q and which qudits are acted upon at each step in the circuit.

The statistical-mechanical interpretation has been a useful one for similar problems in the past, where certain RQC moment quantities can be exactly rewritten as the partition sum over spin configurations of a lattice model, as depicted in the central diagram in Fig. 4. We can arrive at the formulation as in Eq. (B18) from the lattice model by summing over a subset of the spins and reinterpreting the resulting nodes as four-body interaction vertices.

This exact rewriting of RQC moment quantities has been used to compute, for instance, correlation functions [3], Rényi entropies [40], and the distance to forming an approximate design [17]. Moreover, thermal phase transitions in the classical model can be related to phase transitions of entanglement-entropy-like quantities for the output state of the RQC [41–43]. The interpretation is particularly intriguing when considering analogous quantities to Z for higher moments. The collision probability is a second-moment quantity, and the resulting statistical-mechanical model, has Ising-like variables with two possible values. Quantities related to the kth moment will map to classical statistical-mechanical models that have $k!$ possible values, one for each element of the symmetric group S_k. However, one challenge of computing higher-moment quantities is that the weights in the partition function can be negative (corresponding to nonreal values of the energy for certain spin configurations), complicating many strategies for bounding its behavior, including the strategies employed in the rest of this paper.

3. Unbiased random walk

We can build from the formula for Z in Eq. (B18) and reexpress it in terms of a length-s unbiased random walk through configuration space $\{I, S\}^n$, which we denote

P_u. At time step 0, a configuration $\vec{\gamma}^{(0)}$ is chosen uniformly at random, i.e., the initial distribution is the uniform distribution in configuration space, denoted Λ_u. Then configuration $\vec{\gamma}^{(t+1)}$ at time step $t+1$ is generated from the configuration $\vec{\gamma}^{(t)}$ at time step t as follows: letting $A^{(t)} = \{a, b\}$, if the ath and bth bits of $\vec{\gamma}^{(t)}$ agree, then the configuration is left unchanged at time step $t+1$; if they disagree, either the value at a or the value at b is flipped each with probability $1/2$ to form $\vec{\gamma}^{(t+1)}$. The weight is reduced each time a bit is flipped. Thus we can write

$$Z = \frac{2^n}{(q+1)^n} \mathbb{E}_{P_u, \Lambda_u} \left[\left(\frac{2q}{q^2+1} \right)^{\text{(number of bit flips during walk)}} \right],$$
(B19)

where $\mathbb{E}_{P_u, \Lambda_u}$ indicates the expectation over the choosing a length-s walk as described above, where the initial distribution is Λ_u. This is seen to be equivalent to Eq. (B18), since the probability of a certain trajectory occurring is given by $q^{-n}(1/2)^{\text{number of bit flips}}$ and thus each trajectory contributes exactly the same amount toward Z, once the probability of observing the trajectory is accounted for.

4. Biased random walk

A potential problem with the unbiased-random-walk picture is that the weight of a particular walk is related to the number of bit flips that occur during that walk; it depends not only where the walk begins and ends but also on how it got there. To fix this issue, we can form an equivalent *biased* random walk denoted P_b. In this case, the initial distribution Λ_b is not uniform over $\{I, S\}^n$; rather, the probability of choosing $\vec{\gamma}^{(0)} = \vec{v}$ is proportional to $q^{-|\vec{v}|}$, where $|\vec{v}|$ is the Hamming weight of $\vec{v}$ (number of S entries). Specifically, we have

$$\Lambda_b(\vec{v}) = \frac{q^n}{(q+1)^n} q^{-|\vec{v}|}.$$
(B20)

The dynamics of P_b are the same as P_u, except that when the two bits involved in a gate disagree, it chooses to flip the S to I with probability $q^2/(q^2+1)$ and I to S with probability $1/(q^2+1)$. Thus, it is biased in the I direction. Then, we can express

$$Z = \frac{1}{q^n} \mathbb{E}_{P_b, \Lambda_b} \left[q^{|\vec{\gamma}^{(s)}|} \right].$$
(B21)

Note that the quantity being averaged is exponentially large in the Hamming weight of its final ending point, making the quantity sensitive to the probability that the biased walk stays far from the all I configuration. The biased walk is observed to be equivalent to the unbiased walk simply by noting that once the probability of observing a certain trajectory is included, every trajectory contributes the

same amount to Z for both walks. The exponential weighting underneath the expectation in the biased walk exactly cancels the bias in the probability of observing a certain walk.

5. Computing sums over trajectories

Throughout our analysis, we need to compute weighted sums over various trajectories or, relatedly, compute probabilities that the biased and unbiased walks end in a certain place. We use the following lemma. The key takeaway is that (perhaps surprisingly), in the limit of infinite size, the contribution of all trajectories originating from a certain initial configuration *depends only on the Hamming weight* of that initial configuration and not on the configuration itself. Moreover, this contribution can be calculated. This lemma is a more precise and generalized version of the recursive calculation of $Q(x)$ in Sec. V in the main text.

Lemma 1. *Fix an infinite-size circuit diagram, that is, an infinite sequence of qudit pairs $A = (A^{(1)}, A^{(2)}, \ldots)$. Also fix integers $0 \le x, y, m \le n$ such that $y \le x < y + m$, as well as an initial configuration $\vec{\gamma}^{(0)}$ such that $|\vec{\gamma}^{(0)}| = x$. For each $s \ge 0$, let $\mathcal{T}_s$ be the set of length-s trajectories that*

(1) begin at configuration $\vec{\gamma}^{(0)}$
(2) have a nonzero contribution to Z for the circuit diagram $(A^{(1)}, \ldots, A^{(s)})$ formed by truncating A to length s
(3) end at any configuration $\vec{\gamma}^{(s)}$ for which $|\vec{\gamma}^{(s)}| = y$, and
(4) satisfy $y < |\vec{\gamma}^{(t)}| < y + m$ for all $t = 0, 1, 2, \ldots, s - 1$

Let $\mathcal{T} = \bigcup_{s=0}^{\infty} \mathcal{T}_s$. Then,

$$\sum_{\gamma \in \mathcal{T}} \left(\frac{q}{q^2+1} \right)^{\text{(number of bit flips during } \gamma)}$$

$$= \frac{1}{1 - q^{-2m}} \left(q^{-(x-y)} - q^{-2m+x-y} \right).$$
(B22)

Proof. First, we claim that the sum should depend only on x, y, and m and not on $\vec{\gamma}^{(0)}$ (other than through its dependence on x). To see this, note that there is a one-to-one correspondence between trajectories in $\mathcal{T}$ and sequences of Hamming weights $(x, x_1, \ldots, x_{s'-1}, y)$ with the property that either $x_t = x_{t+1} + 1$ or $x_t = x_{t+1} - 1$ for every t (no consecutive duplicates). This is seen by (1) the fact that given a trajectory in $\mathcal{T}$, one can generate such a sequence by taking the Hamming weight of each configuration in the sequence and removing consecutive duplicates and (2) the fact that given such a Hamming weight sequence, one can generate a unique trajectory by starting with $\vec{\gamma}^{(0)}$, evolving the trajectory according to the circuit diagram A, and

always choosing whether to flip I to S or S to I so that the order of Hamming weights prescribed by the sequence is followed. Thus, the sum over trajectories in $\mathcal{T}$ may be replaced by a sum over Hamming weight sequences, which does not depend on $\vec{\gamma}^{(0)}$, except through its Hamming weight x.

For each x in the interval $[y, y + m]$, let the expression on the left-hand side of the lemma be given by $Q(x)$. Then for each x in $[y + 1, y + m - 1]$, we have the recursion relation

$$Q(x) = \frac{q}{q^2 + 1}[Q(x - 1) + Q(x + 1)], \qquad \text{(B23)}$$

since the first bit flip will either send x to $x - 1$ or to $x + 1$ and in either case a factor of $q/(q^2 + 1)$ is incurred. The recursion relation gives rise to a general solution of the form

$$Q(x) = Fq^x + Gq^{-x} \qquad \text{(B24)}$$

for some constants F and G. This is a unique solution, since all values can be generated once two consecutive values are specified, and the specification of two consecutive values also uniquely specifies F and G. To find F and G in this case, we must also impose the boundary conditions $Q(y) = 1$ and $Q(y + m) = 0$, since if $x = y$ the only trajectory in $\mathcal{T}$ is the length-0 trajectory $(\vec{\gamma}^{(0)})$ and, if $x = y + m$, $\mathcal{T}$ is the empty set. By specifying these boundary conditions, we can solve for F and G and verify the statement of the lemma. ∎

Corollary 1. *Fix non-negative integers x, y, m such that $y \leq x < y + m$. For the biased walk, if the starting configuration has Hamming weight x, the probability that the walk reaches a configuration with Hamming weight y before it reaches a configuration with Hamming weight $y + m$ is given by*

$$\frac{q^{x-y}}{1 - q^{-2m}}\left[q^{-(x-y)} - q^{-2m+x-y}\right]. \qquad \text{(B25)}$$

Proof. The transition rules of the biased walk prescribe that transitions upward in Hamming weight occur with probability $1/(q^2 + 1)$ and transitions downward in Hamming weight occur with probability $q^2/(q^2 + 1)$. Thus the probability of a series of transitions in which the initial Hamming weight is x, the final Hamming weight is y, and the number of times a bit flip occurs is b is precisely $q^{x-y}[q/(q^2 + 1)]^b$. The sum over all paths weighted by their probability is then precisely the sum on the left-hand side of Lemma 1 scaled by q^{x-y}, yielding the corollary. ∎

Corollary 2. *If we begin at a trajectory $\vec{\gamma}^{(0)}$ with $|\vec{\gamma}^{(0)}| = x$ and allow the biased walk to evolve until it ends at one of the fixed points I^n or S^n, then the probability that the trajectory ends at I^n is given by*

$$P_I(x) = \frac{1}{1 - q^{-2n}}\left(1 - q^{-2n+2x}\right) \qquad \text{(B26)}$$

and the probability that it ends at S^n is given by

$$P_S(x) = \frac{q^{-2n+2x}}{1 - q^{-2n}}\left(1 - q^{-2x}\right). \qquad \text{(B27)}$$

Proof. Termination at I^n corresponds to the cases where Hamming weight 0 is hit before Hamming weight n. Thus the equation for $P_I(x)$ follows from Corollary 1 with $y = 0$ and $m = n$. We have $P_S(x) = 1 - P_I(x)$, since the trajectory must terminate at one fixed point or the other. ∎

6. Sanity check: Infinite-circuit-size convergence to Haar value

The Markov chain has two stationary distributions, at configurations I^n and S^n. In the infinite-circuit-size limit, the biased walk will converge to a mixture of these two fixed-point configurations, where the amount of mass at each fixed point depends only on the Hamming weight of the initial configuration, as described in Corollary 2. Using the expressions for P_I and P_S, we find that, in the infinite-circuit-size limit,

$$
\begin{aligned}
Z &= \frac{1}{q^n}\sum_{\vec{\gamma}^{(0)}}\Lambda_b(\vec{\gamma}^{(0)})\underset{P_b,\vec{\gamma}^{(0)}}{\mathbb{E}}\left[q^{|\vec{\gamma}^{(s)}|}\right] \\
&= \frac{1}{(q+1)^n}\sum_{\vec{\gamma}^{(0)}}q^{-|\vec{\gamma}^{(0)}|}[P_I(|\vec{\gamma}^{(0)}|) + q^n P_S(|\vec{\gamma}^{(0)}|)] \\
&= \frac{1}{(q+1)^n(1 - q^{-2n})}\sum_{x=0}^{n}\binom{n}{x} \\
&\quad \times q^{-x}(1 - q^{-2n+2x} + q^{-n+2x} - q^{-n}) \\
&= \frac{1}{(q+1)^n(1 - q^{-2n})}\left[\left(\frac{q+1}{q}\right)^n - q^{-2n}(q+1)^n \right. \\
&\quad \left. + q^{-n}(q+1)^n - q^{-n}\left(\frac{q+1}{q}\right)^n\right] \\
&= \frac{(2q^{-n} - 2q^{-2n})(q+1)^n}{(q+1)^n(1 - q^{-2n})} = \frac{2}{q^n + 1} = Z_H, \quad \text{(B28)}
\end{aligned}
$$

where Z_H is the Haar value. This outcome is expected, since in the infinite-circuit-size limit, the distribution over random unitaries formed from Haar-random local components will approach the distribution over n-qudit Haar random unitaries.

APPENDIX C: BOUNDS FOR GENERAL ARCHITECTURES

1. Upper bound on collision probability

In order to have a meaningful upper bound, we need the architecture to satisfy basic connectivity requirements; for example, if the architecture performs gates on the same pair of qudits over and over again, Z will never decrease and the output distribution will never become anticoncentrated. We need to rule out this sort of architecture.

Recall that a RQC architecture is a (possibly randomized) procedure for choosing a length-s sequence $(A^{(1)}, \ldots, A^{(s)})$ of pairs of qudit indices on which to perform a Haar-random gate.

Definition 5 (Regularly connected)**.** *We say that a RQC architecture is r-regularly connected if, for any n, any t, any subsequence $A = (A^{(1)}, \ldots, A^{(t)})$, and any proper subset $R \subset [n]$ of qudit indices, there is at least a $1/2$ probability that, conditioned on the first t gates in the gate sequence being A, there exists some index t' for which $t < t' \le t + rn$, $A^{(t')} \cap R \ne \emptyset$, and $A^{(t')} \not\subset R$.*

The above definition requires that given any partition of the qudits into two sets, we should expect at least one gate to couple a qudit from one set with a qudit from the other set after only a linear number of gates. Note that both the 1D and the complete-graph architecture have this property. In 1D, it only takes two layers, or n gates, to guarantee having performed a gate that crosses any partition one might choose. Similarly, in the complete-graph architecture, the probability that a randomly chosen gate crosses a partition is at least $1/n$ (which happens if the partition splits the indices into a set with one index and a set with the other $n - 1$ indices) and the probability of having crossed the partition becomes large after $\Theta(n)$ gates. Most natural architectures that we might consider have this property. One architecture that is not regularly connected is the hypercube architecture, where $n = 2^D$ qudits lie at the vertices of a D-dimensional hypercube and D layers of gates are performed by cycling through each set of parallel edges. In this architecture, it would take $nD/2 = \Theta(n \log(n))$ gates to guarantee that any partition has been crossed.

Assuming the regularly connected property, we can show a weak upper bound on the collision probability.

Theorem 7 (Theorem 3 from main text)**.** *If a RQC architecture is r-regularly connected, then the collision probability satisfies*

$$Z \le Z_H \left[1 + e^{-\frac{2a}{n}(s - s^*)} \right], \qquad (C1)$$

where

$$a := (2r)^{-1} \log \left[\frac{2(q^2 + 1)}{(q + 1)^2} \right], \qquad (C2)$$

$$s^* := (2a)^{-1} \log \left(\frac{2q}{q + 1} \right) n^2 + O(n). \qquad (C3)$$

Proof. We use the expression given to us by the unbiased walk in Eq. (B19):

$$Z = \frac{2^n}{(q + 1)^n} \mathop{\mathbb{E}}_{P_u, \Lambda_u} \left[\left(\frac{2q}{q^2 + 1} \right)^{\text{(number of bit flips during walk)}} \right]. \qquad (C4)$$

Define $Z^{(t)}$ to be the value of the collision probability, given above via the biased walk, after t time steps, so that $Z = Z^{(s)}$ and $Z^{(0)} = 2^n/(q + 1)^n$.

Consider a given trajectory produced by the unbiased walk up to time step t, $\gamma = (\vec{\gamma}^{(0)}, \ldots, \vec{\gamma}^{(t)})$. If $\vec{\gamma}^{(t)} = I^n$ or $\vec{\gamma}^{(t)} = S^n$, then the walk has reached a fixed point and will never again change. From the calculation in Sec. 6, we know that the sum of all the weights of all walks *of any length* that reach a fixed point is precisely Z_H. Since the weights are non-negative, this implies that the sum over walks that have reached it before time step t is less than Z_H and hence the combined weight of trajectories that have *not* reached a fixed point by time step t is at least $Z^{(t)} - Z_H$. Meanwhile, if $\vec{\gamma}^{(t)}$ is not at a fixed point, then we can consider the proper subset $R \subset [n]$ of sites with value S. By the r-regularly connected property, there is at least a $1/2$ chance that one of the gates between time step $t + 1$ and $t + rn$ matches an index in R with one in the complement of R. When this happens, a bit must be flipped and the weight of that trajectory is reduced by factor $2q/(q^2 + 1)$. Thus, the following must hold:

$$Z^{(t+rn)} - Z_H \le \left(\frac{1}{2} + \frac{1}{2} \frac{2q}{q^2 + 1} \right) \left[Z^{(t)} - Z_H \right]. \qquad (C5)$$

Moreover, we know that $Z^{(0)} = 2^n/(q + 1)^n$, so

$$
\begin{aligned}
Z^{(s)} &\le Z_H + \left[\frac{(q + 1)^2}{2(q^2 + 1)} \right]^{s/(rn)} \left[\frac{2^n}{(q + 1)^n} - Z_H \right] \\
&\le Z_H + \left[\frac{(q + 1)^2}{2(q^2 + 1)} \right]^{s/(rn)} \left[\frac{2^n}{(q + 1)^n} \right] \\
&= Z_H \left\{ 1 + \frac{2^n(q^n + 1)}{2(q + 1)^n} \left[\frac{(q + 1)^2}{2(q^2 + 1)} \right]^{s/(rn)} \right\} \\
&\le Z_H \left\{ 1 + \frac{2^n q^n}{(q + 1)^n} \left[\frac{(q + 1)^2}{2(q^2 + 1)} \right]^{s/(rn)} \right\} \\
&\le Z_H [1 + e^{-\frac{2a}{n}(s - s^*)}], \qquad (C6)
\end{aligned}
$$

where

$$a := (2r)^{-1} \log \left[\frac{2(q^2+1)}{(q+1)^2} \right] = \Theta(1), \qquad \text{(C7)}$$

$$s^* := (2a)^{-1} \log \left(\frac{2q}{q+1} \right) n^2 = \Theta(n^2). \qquad \text{(C8)}$$

∎

Note that we make no attempt to optimize the constant prefactor of the $\Theta(n^2)$ or the value of a. Indeed, we conjecture that Theorem 7 could be improved so that $s^* = \Theta(n \log(n))$, which would be a dramatic improvement that implies that the fundamental scaling of the anticoncentration size is independent of the connectivity of the architecture, so long as it satisfies the regularly connected property.

2. Lower bound on collision probability

In this section, we prove an $\Omega(n \log(n))$ lower bound on the circuit size needed for anticoncentration in general circuit architectures. This also implies an $\Omega(\log(n))$ lower bound on the anticoncentration depth.

Theorem 8 (Theorem 4 from main text). *For any RQC architecture of size s on n qudits with local dimension q, the collision probability satisfies*

$$Z \geq \frac{Z_H}{2} \exp \left\{ \frac{\log(q)}{q+1} \exp \left[\log(n) - \frac{2s}{n} \log(q^2+1) \right] \right\}. \qquad \text{(C9)}$$

Corollary 3. *For a given RQC architecture, let s_{AC} be the minimum circuit size, as a function of n, such that $Z \leq 2Z_H$. Then, it must hold that*

$$s_{\mathrm{AC}} \geq \left[2 \log(q^2+1) \right]^{-1} n \log(n) - O(n). \qquad \text{(C10)}$$

Proof. This statement follows directly from the bound in Eq. (C9). ∎

Corollary 4. *For a given RQC architecture, let d_{AC} be the minimum circuit depth, as a function of n, such that $Z \leq 2Z_H$. Then, it must hold that*

$$d_{\mathrm{AC}} \geq \left[\log(q^2+1) \right]^{-1} \log(n) - O(1). \qquad \text{(C11)}$$

Proof. Each layer can have at most $n/2$ gates, so it must hold that $d_{\mathrm{AC}} \geq 2s_{\mathrm{AC}}/n$. ∎

Proof of Theorem 8. We use the framework of the biased random walk, given by the expression for Z in Eq. (B21). For each of the n sites, there is some initial probability that

it starts with value S and then each gate involving that site has some chance of flipping it to value I. However, there will always be some minimum probability that even after many gates, the value has not yet been flipped to I. This constitutes the idea behind our lower bound.

Given an index $j \in [n]$, we compute a lower bound on the probability that $\gamma_j^{(t)} = S$ for all $t = 0, 1, \ldots, s$ (i.e., the jth bit begins with value S and is never flipped to I), as a function of the number of gates s_j that act on qudit j,

$$\Pr_{P_b, \Lambda_b} [\gamma_j^{(t)} = S \ \forall t \in \{0, \ldots, s\}] \geq \frac{1}{q+1} \left(\frac{1}{q^2+1} \right)^{s_j}, \qquad \text{(C12)}$$

since there is a $1/(q+1)$ chance that $\gamma_j^{(0)} = S$ when we draw $\vec{\gamma}^{(0)}$ from Λ_b and the probability it does not flip after each gate is at least $1/(q^2+1)$. This holds for each j and thus we have

$$\mathbb{E}_{P_b, \Lambda_b} [|\vec{\gamma}^{(s)}|] = \sum_{j=1}^{n} \Pr_{P_b, \Lambda_b} [\gamma_j^{(s)} = S]$$

$$\geq \sum_{j=1}^{n} \Pr_{P_b, \Lambda_b} [\gamma_j^{(t)} = S \ \forall t \in \{0, \ldots, s\}]$$

$$\geq \frac{1}{q+1} \sum_{j=1}^{n} \left(\frac{1}{q^2+1} \right)^{s_j}. \qquad \text{(C13)}$$

Since each of the s gates in the circuit diagram acts on two indices, it must hold that $\sum_j s_j = 2s$ and, given this constraint, the minimum of the final expression above occurs when all the s_j are equal and thus

$$\mathbb{E}_{P_b, \Lambda_b} [|\vec{\gamma}^{(s)}|] \geq \frac{1}{q+1} n \left(\frac{1}{q^2+1} \right)^{2s/n}. \qquad \text{(C14)}$$

By convexity of the exponential function, we have $\mathbb{E}[q^x] \geq q^{\mathbb{E}[x]}$ and hence

$$Z = \frac{1}{q^n} \mathbb{E}_{P_b, \Lambda_b} [q^{|\vec{\gamma}^{(s)}|}]$$

$$\geq \frac{1}{q^n} \exp \left[\log(q) \frac{n}{q+1} \left(\frac{1}{q^2+1} \right)^{2s/n} \right]$$

$$\geq \frac{Z_H}{2} \exp \left\{ \frac{\log(q)}{q+1} \exp \left[\log(n) - \frac{2s}{n} \log(q^2+1) \right] \right\}. \qquad \text{(C15)}$$

∎

APPENDIX D: BOUNDS FOR THE 1D ARCHITECTURE

We now focus specifically on the 1D architecture defined formally in Definition 2. We assume periodic boundary conditions, although it would be possible to consider open boundary conditions as well. In 1D, the qudits are arranged in a geometrically local fashion and it is fruitful to think of a configuration $\vec{\gamma} \in \{I, S\}^n$ as being composed of contiguous *domains*, consecutive sites where all the values are I or all the values are S. We then identify *domain walls* as locations where one domain ends and another begins. Gates that couple qudits in different domains cause one of the values to flip, which moves the domain wall separating those domains one unit to the left or one unit to the right. The notation for talking formally about this is discussed in the next subsection and then the upper and lower bounds on Z are proved.

1. Domain walls and notation

In 1D, configurations $\vec{\gamma} \in \{I, S\}^n$ are associated with a set of domain-wall locations. We let

$$DW(\vec{\gamma}) := \{e \in \{0, 1, 2, \ldots, n-1\} : \gamma_e \neq \gamma_{e+1}\} \quad \text{(D1)}$$

be the set of domain-wall positions for a configuration $\vec{\gamma}$, where γ_0 is identified with γ_n when there are periodic boundary conditions. For each set of domain-wall locations, there are exactly two configurations that map to it, since choosing $\gamma_0 = I$ or $\gamma_0 = S$ determines the value of all other sites.

A configuration trajectory $\gamma = (\vec{\gamma}^{(0)}, \ldots, \vec{\gamma}^{(s)})$ is then associated with a sequence of sets of domain-wall locations $G = (g^{(0)}, \ldots, g^{(s)})$, where $g^{(t)} = DW(\vec{\gamma}^{(t)})$. We call G a *domain-wall trajectory*. Domain-wall trajectories with nonzero contribution to the collision probability Z obey the following rules: when there is a domain wall at position e and a gate acts on qudits $\{e, e+1\}$, the domain wall must move to position $e-1$ or $e+1$ (at the cost of a reduction in the weight) and may annihilate with another domain wall if there is already a domain wall at the new position. However, pairs of domain walls cannot be created; the number of domain walls that exist throughout the domain-wall trajectory is nonincreasing and a particular domain wall can be uniquely tracked throughout each step of the trajectory (either until the final step or until its annihilation). Let $\mathcal{G}$ be the set of all domain-wall trajectories that obey these rules. Any domain-wall trajectory $G \in \mathcal{G}$ will have the property that when t is odd, e is even for all $e \in g^{(t)}$, and when t is even (but nonzero), e is odd for all $e \in g^{(t)}$. This is because odd- (even-) numbered layers couple qudits $\{2j-1, 2j\}$ ($\{2j, 2j+1\}$), meaning that domain walls must lie between qudit positions $2j$ and $2j+1$ (between qudit positions $2j-1$ and $2j$) for some j.

By converting the sum over trajectories in Eq. (B18) to a sum over domain-wall trajectories, we can express Z by the equation

$$Z = \frac{2}{(q+1)^n} \sum_{G \in \mathcal{G}} \text{weight}(G), \quad \text{(D2)}$$

where the weight is given as follows, recalling that $A^{(t)}$ is the pair of qudit indices involved in the tth gate, which in 1D is always $A^{(t)} = \{j, j+1\}$ for some j:

$$\text{weight}(G) := \prod_{t=1}^{s} M^{(t)}_{g^{(t-1)} g^{(t)}}, \quad \text{(D3)}$$

$$M^{(t)}_{g^{(t-1)} g^{(t)}} := \begin{cases} \dfrac{q}{q^2+1}, & \text{if } \min(A^{(t)}) \in g^{(t-1)}, \\ 1, & \text{otherwise.} \end{cases} \quad \text{(D4)}$$

In other words, if the gate on qudits $\{j, j+1\}$ and there is a domain wall at position j, then the weight is reduced by a factor $q/(q^2+1)$ (and the domain wall must move to position $j-1$ or position $j+1$, possibly annihilating if a domain wall already exists at that position).

Given two domain-wall trajectories G and G', we consider the combined domain-wall trajectory

$$G \sqcup G' := (g^{(0)} \sqcup g'^{(0)}, \ldots, g^{(s)} \sqcup g'^{(s)}), \quad \text{(D5)}$$

where $\sqcup$ is the disjoint union and is defined only under the assumption $g^{(t)} \cap g'^{(t)} = \emptyset$ for all t.

The upshot of thinking about trajectories in this way is that if $H = G \sqcup G'$, then

$$\text{weight}(H) = \text{weight}(G) \, \text{weight}(G'). \quad \text{(D6)}$$

In particular, we find it useful to decompose a domain-wall trajectory G into $G = G_U \sqcup G_0$, where G_U is a domain-wall trajectory with a conserved number of domain walls throughout the trajectory and G_0 is a trajectory for which $|G_0^{(s)}| = 0$, i.e., all the domain walls have annihilated by the end of the trajectory. This decomposition is unique and an example is shown in Fig. 5. Let $\mathcal{G}_U$ and $\mathcal{G}_0$ be the subsets of $\mathcal{G}$ that have no annihilations and that have no surviving domain walls at the end of the circuit, respectively. Let $\mathcal{G}_{U,k}$ be the subset of $\mathcal{G}_U$ with k domain walls. When the boundary conditions are periodic, k must be even for $\mathcal{G}_{U,k}$ to be nonempty.

2. Collision probability upper bound

Theorem 9 (Theorem 1 from main text). *For the 1D architecture, let*

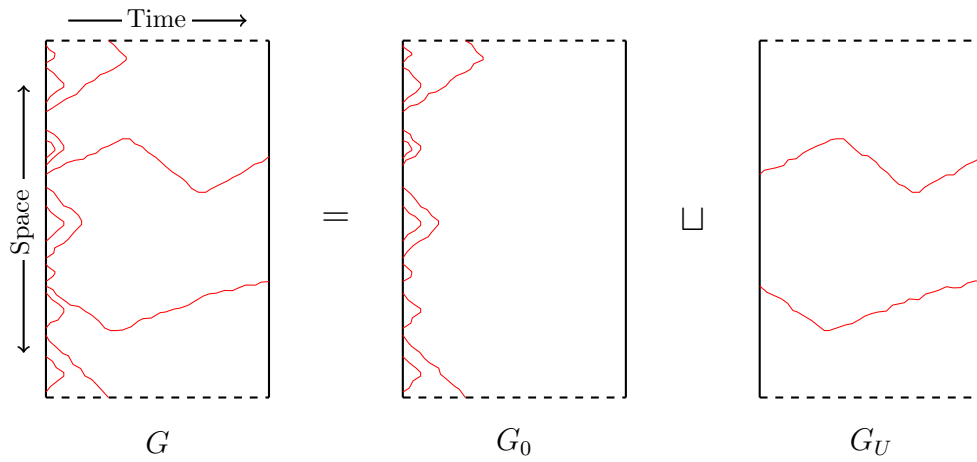$$a := \log\left(\frac{q^2+1}{2q}\right), \quad \text{(D7)}$$

FIG. 5. A sketch illustrating unique decomposition of a domain-wall trajectory G into a disjoint union of one part, G_0, where all domain walls annihilate prior to the end of the circuit, and another part, G_U, where no domain walls annihilate.

$$s^* := \frac{1}{2a} n \log(n) + n \left[\frac{1}{2a} \log(e-1) + \frac{1}{2} \right]$$
$$= (2a)^{-1} n \log(n) + O(n). \tag{D8}$$

Then,

$$Z \leq Z_H [1 + e^{-\frac{2a}{n}(s-s^*)}] \tag{D9}$$

whenever $s \geq s^*$. The circuit depth d is $d = 2s/n$, so we may define $d^* = 2s^*/n$ and equivalently conclude that

$$Z \leq Z_H [1 + e^{-a(d-d^*)}]. \tag{D10}$$

Note that when $s < s^*$, an upper bound on Z can still be inferred from this method. The essence of the proof of Theorem 9 is the same as the proof of the statement proved in Ref. [32], although we express it here within our notation and framework.

Proof. We use the formula in Eq. (D2), which expresses Z as a weighted sum over domain-wall trajectories. Each domain-wall trajectory $G = (g^{(0)}, \dots, g^{(s)})$ can be associated with an integer $k = |g^{(s)}|$, the number of domain walls that remain unannihilated at the end of the trajectory. Due to periodic boundary conditions, k must be even, and we let $k_0 = k/2$. Let $\mathcal{G}_k \subset \mathcal{G}$ be the associated set of length-s domain-wall trajectories and let $\mathcal{G}_{U,k} \subset \mathcal{G}_k$ be the subset containing domain-wall trajectories that have a conserved number of domain walls throughout. As discussed in the previous subsection, it is possible to uniquely decompose $H \in \mathcal{G}_k$ into $H = G \sqcup G'$, where $G \in \mathcal{G}_{U,k}$ and $G' \in \mathcal{G}_0$.

Suppose that we fix a domain-wall configuration $g^{(0)}$ for the initial time step at the beginning of the circuit with k domain walls. There are $\binom{n}{k}$ such configurations. The total weight of all the trajectories in $\mathcal{G}_{U,k}$ that begin at this configuration is at most $[2q/(q^2+1)]^{k(d-1)}$, since each domain wall must move either left or right (introducing a factor of 2) during each of the d layers of gates, except for

possibly the first layer (if the domain wall begins at an even position, it does not move during the first layer), and each time one moves it incurs a weight reduction of $q/(q^2+1)$. This does not account for the rule that the k domain walls cannot intersect but it still yields an upper bound on the total weight.

Meanwhile, the sum of the weights of all domain-wall trajectories in $\mathcal{G}_0$ approaches $Z_H(q+1)^n/2$ from below as the depth increases. This follows from the analysis in Sec. 6, where it is shown that the sum over all trajectories that eventually reach a fixed point is exactly $Z_H(q+1)^n$ but at a finite depth not every trajectory will have reached a fixed point, so only a subset of the terms are included in the sum. Due to the fact that each domain-wall configuration corresponds to two equal-weight trajectories through $\{I, S\}^n$, the sum of the weights of all the domain-wall trajectories in $\mathcal{G}_0$ can be at most $Z_H(q+1)^n/2$.

Collecting these observations, and recalling that $k = 2k_0$, we have

$$Z = \frac{2}{(q+1)^n} \sum_{k_0=0}^{n/4} \sum_{G \in \mathcal{G}_{2k_0}} \text{weight}(G)$$

$$= \frac{2}{(q+1)^n} \sum_{k_0=0}^{n/4} \sum_{\substack{G \in \mathcal{G}_{U,2k_0} \\ G' \in \mathcal{G}_0 \\ G \cap G' = \emptyset}} \text{weight}(G) \, \text{weight}(G')$$

$$\leq \left[\sum_{k_0=0}^{n/4} \sum_{G \in \mathcal{G}_{U,2k_0}} \text{weight}(G) \right] \left[\frac{2}{(q+1)^n} \sum_{G' \in \mathcal{G}_0} \text{weight}(G') \right]$$

$$\leq \left[\sum_{k_0=0}^{n/4} \binom{n}{2k_0} \left(\frac{2q}{q^2+1} \right)^{2k_0(d-1)} \right] (Z_H)$$

$$= Z_H \sum_{k_0=0}^{n/4} \binom{n}{2k_0} (e^{-a})^{2k_0(d-1)}$$

$$\leq Z_H[1 + e^{-a(d-1)}]^n$$

$$\leq Z_H[1 + (e-1)ne^{-a(d-1)}] \tag{D11}$$

$$= Z_H \{1 + \exp[\log(n) - da + \log(e-1) + a]\}$$

$$\leq Z_H \{1 + \exp[-a(d - d^*)]\}, \tag{D12}$$

where Eq. (D11) holds so long as $d \geq d^*$, by application of Lemma 2. ∎

Lemma 2. *If $b, c > 0$ and $cb \leq 1$, then*

$$(1 + c)^b \leq 1 + cb(e - 1). \tag{D13}$$

Proof.

$$(1+c)^b = \sum_{k=0}^{b} \binom{b}{k} c^k = 1 + cb \sum_{k=1}^{b} \binom{b}{k} \frac{c^{k-1}}{b}$$

$$\leq 1 + cb \sum_{k=1}^{b} \binom{b}{k} b^{-k} \leq 1 + cb \left[(1 + b^{-1})^b - 1\right]$$

$$\leq 1 + cb(e - 1). \tag{D14}$$

∎

3. Collision probability lower bound

Theorem 10 (Theorem 5 from main text). *Consider the 1D architecture. There are constants A and A' such that as long as $s^* - s \geq A'n$, the collision probability satisfies*

$$Z \geq \frac{Z_H}{2} \exp\left[Ae^{\log(n) - \frac{2a}{n}s}\right]. \tag{D15}$$

where a and s^* are the same as in Theorem 9.

In our proof, the constant A is explicit but very small, on the order of e^{-10}, and $A' \approx -\log(A)$. The value of A could certainly be improved with some attempt at optimization.

Corollary 5. *For the 1D architecture, if we define s_{AC} and d_{AC} to be the smallest circuit size and circuit depth for which $Z \leq 2Z_H$, then*

$$\left| s_{\mathrm{AC}} - \left[2\log\left(\frac{q^2 + 1}{2q}\right)\right]^{-1} n\log(n) \right| \leq O(n), \tag{D16}$$

$$\left| d_{\mathrm{AC}} - \left[\log\left(\frac{q^2 + 1}{2q}\right)\right]^{-1} \log(n) \right| \leq O(1). \tag{D17}$$

Proof. Theorem 9 implies that

$$s_{\mathrm{AC}} \leq s^* = (2a)^{-1} n\log(n) + O(n). \tag{D18}$$

Meanwhile, Theorem 10 implies that if

$$s \leq (2a)^{-1} n\log(n) - \max\left\{(2a)^{-1}\log[\log(4)A^{-1}], A'\right\}$$

$$n = (2a)^{-1} n\log(n) - O(n), \tag{D19}$$

then $Z \geq 2Z_H$. Hence $s_{\mathrm{AC}} \geq (2a)^{-1} n\log(n) - O(n)$. Together, these imply that $|s_{\mathrm{AC}} - (2a)^{-1} n\log(n)| = O(n)$. ∎

Proof of Theorem 10. Equation (D2) expresses Z as a weighted sum over domain-wall trajectories. Heuristically, when $d < d^*$, we expect that the output distribution will *not* be anticoncentrated and that domain-wall trajectories drawn at random with probability proportional to its weight will usually have many domain walls that never annihilate. To lower bound Z, we sum over the set of configurations with k unannihilated domain walls for a particularly chosen value of k.

For a fixed value of the depth d, define

$$n_H := \frac{e^{(d-1)a}}{2(e-1)}. \tag{D20}$$

We choose n_H to be exactly half the value of n for which a depth-d circuit would be anticoncentrated. Heuristically, we expect on the order of $n/2n_H$ unannihilated domain walls in typical configurations.

Let k be an even integer to be specified later. Let $\mathcal{H}_k \subset \mathcal{G}_{U,k}$ be the set that contains any domain-wall trajectory $H = (h^{(0)}, \ldots, h^{(s)})$ for which:

(1) H has k domain walls at each time step (none annihilate).
(2) For each of the k domain walls in the initial configuration $h^{(0)}$, the nearest domain wall in both directions is at most n_H positions away.

Now, temporarily fix some $H \in \mathcal{H}_k$. It has k domain walls that move around throughout the trajectory. We let $e_{H,j,t}$ be the location of the jth domain wall at time step t in the trajectory H. We then define the set $\mathcal{J}_{H,j} \subset \mathcal{G}_0$, for $j = 1, \ldots, k$ to be the set of domain-wall trajectories for which (1) all of the domain walls annihilate before time step s and (2) the position e_t of any domain wall at time step t satisfies

$$e_{H,j,t} < e_t < e_{H,j+1,t}. \tag{D21}$$

In other words, all of the domain walls fall between the jth and $(j + 1)$th domain walls of H. This ensures that H is disjoint from any $J_j \in \mathcal{J}_{H,j}$.
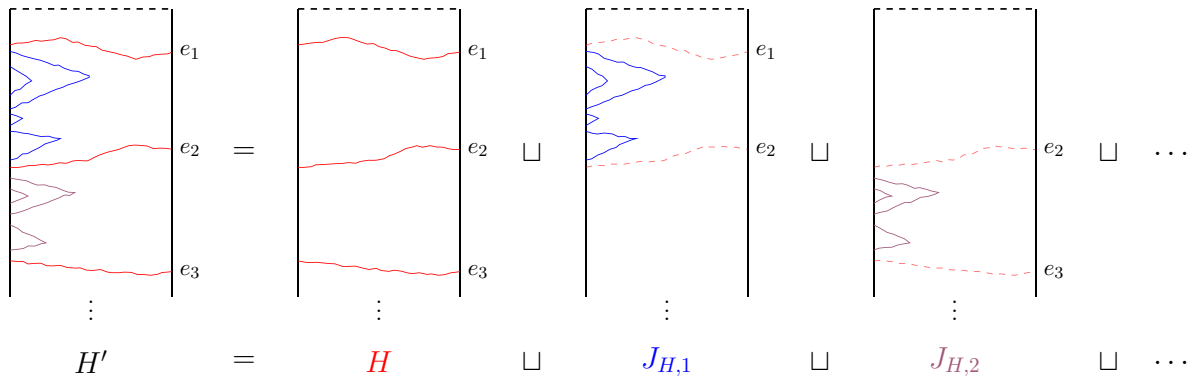
FIG. 6.   The outline of the main idea of the proof of Theorem 10. We choose a domain-wall trajectory H that has k domain walls that never annihilate and such that the distance between consecutive domain walls is always at most n_H. We then choose domain-wall trajectories $J_{H,1}, \ldots, J_{H,k}$ such that the domain walls of $J_{H,j}$ lie between the jth and $(j+1)$th domain walls of H and all annihilate before the end of the circuit. The domain-wall configuration H' is the disjoint union of H and $J_{H,j}$ for $j = 1, \ldots, k$. We can lower bound the collision probability by lower bounding the weighted sum over the contribution from all H' formed this way.

The specification of a trajectory $H \in \mathcal{H}_k$ as well as $J_j \in \mathcal{J}_{H,j}$, for each $j = 1, \ldots, k$, determines a unique trajectory $H' = H \sqcup J_{H,1} \sqcup \ldots \sqcup J_{H,k}$. This decomposition is illustrated in Fig. 6. Thus, if we perform the weighted sum only over the set of H' formed in this way, we arrive at a lower bound to Z, as follows:

$$Z = \frac{2}{(q+1)^n} \sum_{H \in \mathcal{G}} \text{weight}(H)$$

$$\geq \frac{2}{(q+1)^n} \left[\sum_{H \in \mathcal{H}_k} \text{weight}(H) \right] \left[\sum_{J_1 \in \mathcal{J}_{H,1}} \text{weight}(J_1) \right]$$

$$\cdots \left[\sum_{J_k \in \mathcal{J}_{H,k}} \text{weight}(J_k) \right]. \tag{D22}$$

The quantities in parentheses can be bounded with the following two lemmas, the proofs of which are delayed until after the proof of the theorem.

Lemma 3. *If $4d \leq \lfloor n/k \rfloor$ and $n_H/2 \geq \lceil n/k \rceil$ hold, then the set $\mathcal{H}_k$ satisfies*

$$\sum_{H \in \mathcal{H}_k} \text{weight}(H) \geq \left(\frac{1}{2} \left\lfloor \frac{n}{k} \right\rfloor \right)^k \left(\frac{2q}{q^2+1} \right)^{dk}. \tag{D23}$$

Lemma 4. *Fix a value of H and j. Suppose that the jth and $(j+1)$th domain walls of the initial configuration of H lie at positions e and $e + X - 1$ (mod n), respectively, for some positive integer $X < n$. Then,*

$$\left[\sum_{J \in \mathcal{J}_{H,j}} \text{weight}(J) \right] \geq \frac{1}{c} \left(\frac{q+1}{q} \right)^X, \tag{D24}$$

where $c = 3e^{10}$.

The sum of the domain length X for each of the domains is simply n. Thus the $((q+1)/q)^X$ factors cancel the $1/(q+1)^n$ prefactor for Z, and we have

$$Z \geq q^{-n} \left(\frac{1}{2} \left\lfloor \frac{n}{k} \right\rfloor \right)^k c^{-k} \left(\frac{2q}{q^2+1} \right)^{dk}$$

$$= q^{-n} \left(\frac{1}{2} \left\lfloor \frac{n}{k} \right\rfloor \right)^k c^{-k} e^{-adk} \tag{D25}$$

for any k that satisfies $4d \leq \lfloor n/k \rfloor$ and $n_H/2 \leq \lceil n/k \rceil$.

Now, we choose a value of k to maximize the right-hand side of the above equation. In the limit of large n, the requirement that k is an even integer will have negligible effect. In our analysis, we handle this requirement by defining k' to be a real number and k to be the smallest even integer larger than k' and then we make a few rather crude bounds on the floor and ceiling of quantities such as n/k, which are not asymptotically tight but are good enough for our purposes. We choose

$$k' := \frac{n \left(\frac{2q}{q^2+1} \right)^d}{8ce} = \frac{ne^{-da}}{8ce} = \frac{n}{n_H} \frac{e^{-a}}{16e(e-1)c}, \tag{D26}$$

$$k := \text{smallest even integer greater than } k'. \tag{D27}$$

Note that n/k' is at least $8ce$, which is very large, meaning that $\lceil n/2k' \rceil/2 \leq n/2k \leq 2\lfloor n/2k' \rfloor$ certainly holds. For finite n, we can say that as long as $k' \geq 1$, then $k' \leq k \leq 2k'$ will hold. The requirement $k' \geq 1$ translates into

$$d \leq a^{-1}[\log(n) - \log(8ce)], \tag{D28}$$

which, by recalling $s = nd/2$ and that $s^* \geq (2a)^{-1} n \log(n)$, can be reexpressed as

$$s^* - s \geq A'n, \tag{D29}$$

with $A' := (2a)^{-1} \log[8ce(e-1)] + 2^{-1}$, which is assumed to hold in the theorem statement. This implies that

$$\left\lfloor \frac{n}{k} \right\rfloor \geq \left\lfloor \frac{n}{2k'} \right\rfloor \geq \frac{n}{4k'}. \tag{D30}$$

Inspection of the formula for k' reveals that the relation $4d \leq \lfloor n/k \rfloor$ holds for any d and n. Moreover, we have $n_H/2 = (ne^{-a})/[k' 32e(e-1)c] \leq \lceil n/k \rceil$, so the second relation holds as well.

Recall that $Z_H = 2/(q^n + 1) \leq 2q^{-n}$. If we substitute the above bound on $\lfloor n/k \rfloor$ into Eq. (D25), we find that

$$\begin{aligned}
Z &\geq \frac{Z_H}{2} \exp(k) \geq \frac{Z_H}{2} \exp(k') \\
&\geq \frac{Z_H}{2} \exp\left(\frac{ne^{-da}}{8ce}\right) \\
&= \frac{Z_H}{2} \exp\left[\frac{1}{8ce} e^{\log(n) - da}\right] \\
&= \frac{Z_H}{2} \exp\left[\frac{1}{8ce} e^{\log(n) - \frac{2as}{n}}\right] \\
&= \frac{Z_H}{2} \exp\left[A e^{\log(n) - \frac{2as}{n}}\right],
\end{aligned} \tag{D31}$$

for $A := 1/8ce$. Note that this value of A is quite small (on the order of e^{-10}) but with some optimization could likely be made much larger. ∎

Now, we provide the delayed proofs of the two lemmas.

Proof of Lemma 3. Each term in the sum on the left-hand side is non-negative, so we make a lower bound by summing over a subset of the terms. To do so, we can split the n indices up into k nearly equal-size segments of length at most $\lceil n/k \rceil$, which is less than $n_H/2$ by assumption. Then, for each of these segments, we choose the location of a single domain wall that is at least distance d from each edge of the segment. This will generate a unique initial domain-wall configuration that satisfies criteria (2) of $\mathcal{H}_k$, since any pair of consecutive domain walls is closer than n_H apart. The total number of choices is at least

$$\left(\left\lfloor \frac{n}{k} \right\rfloor - 2d\right)^k, \tag{D32}$$

which, by the assumption $4d \leq \lfloor n/k \rfloor$, is at least $(\lfloor n/k \rfloor/2)^k$.

Once the initial k domain-wall locations have been chosen, we examine how they can propagate through the circuit. Each layer of gates will force each of the k domain

walls to move in one of two directions and the weight is reduced by a factor $[q/(q^2 + 1)]^k$, except for the first layer, where some of the domain walls may not move if they begin at an even index. Since, by construction, there are no instances where domain walls start within a distance of $2d$ of any other domain wall, there is no chance of domain walls crossing. Thus, we find that for each initial set of k locations chosen in the manner outlined above, the combined weight of all possible trajectories is at least $[2q/(q^2 + 1)]^{kd}$. This proves the lemma. ∎

Proof of Lemma 4. Consider an alternative 1D qudit system with periodic boundary conditions consisting of X sites by identifying site $e + X$ with site e and ignoring all other sites. Because $H \in \mathcal{H}_k$, we can be assured that $X \leq n_H$. Let $\mathcal{J}'_{H,j}$ be the set of all domain-wall trajectories on the size-X system. Let $\mathcal{J}'_{H,j,l}$ be the subset that have $l = 2l_0$ domain walls on the last time step. Because the collision probability, denoted Z_X, for this X-qudit system must satisfy $Z_X \geq Z_{H,X}$, and here $Z_{H,X} = 2/(q^X + 1)$, it must be the case that

$$\begin{aligned}
Z_X &:= \frac{2}{(q+1)^X} \left[\sum_{J' \in \mathcal{J}'_{H,j}} \text{weight}(J)\right] \\
&= \frac{2}{(q+1)^X} \sum_{l_0=0}^{X/4} \left[\sum_{J' \in \mathcal{J}'_{H,j,2l_0}} \text{weight}(J)\right] \\
&\geq \left(\frac{2}{q^X + 1}\right) =: Z_{H,X}.
\end{aligned} \tag{D33}$$

We can upper bound the contribution of all the terms with $l_0 > 0$ in the above expression by the method that yields the upper bound in Theorem 9. The sum of those terms is upper bounded by the second term in Eq. (D11), that is,

$$\begin{aligned}
\frac{2}{(q+1)^X} &\sum_{l_0=1}^{X/4} \left[\sum_{J' \in \mathcal{J}'_{H,j,2l_0}} \text{weight}(J)\right] \\
&\leq Z_{H,X}(e-1) X e^{-a(d-1)} \\
&= \left(\frac{2}{q^X + 1}\right) \frac{X}{2n_H} \leq \frac{1}{2}\left(\frac{2}{q^X + 1}\right),
\end{aligned} \tag{D34}$$

since $X \leq n_H$. Combining Eqs. (D33) and (D34), we find a lower bound on the $l_0 = 0$ term:

$$\sum_{J' \in \mathcal{J}'_{H,j,0}} \text{weight}(J)$$

$$\geq \left[\frac{(q+1)^X}{2}\right]\left(\frac{2}{q^X+1}\right)\left(1-\frac{1}{2}\right)$$

$$= \left(\frac{q+1}{q}\right)^X\left(\frac{1}{2}\frac{q^X}{q^X+1}\right) \geq \left(\frac{q+1}{q}\right)^X\left(\frac{1}{3}\right),$$
$$(D35)$$

where the last inequality follows since $q \geq 2$ and $X \geq 1$ must be true. Now, every domain-wall trajectory in $\mathcal{J}_{H,j}$ will also be in $\mathcal{J}'_{H,j,0}$ but the converse will not be true. Some trajectories in the latter set will have one or more domain walls that intersect with either the jth or the $(j+1)$th domain wall of H at some time step, which is not allowed within the former set. Thus, the sum over the domain-wall trajectories in $\mathcal{J}_{H,j}$ will be smaller than the sum over those in $\mathcal{J}'_{H,j,0}$ but, we argue, by at most some constant factor due to the following argument, which is also described in Fig. 7. Let $\mathcal{B}_{H,j}$ be the set of all trajectories in which every domain wall either intersects the jth domain wall of H at some time step t or it annihilates with a domain wall that previously intersected with the jth domain wall of H. Then, any trajectory in $\mathcal{J}'_{H,j,0}$ can be formed as the disjoint union of a trajectory in $J \in \mathcal{J}_{H,j}$, a trajectory in $K \in \mathcal{B}_{H,j}$, and a trajectory in $K' \in \mathcal{B}_{H,j+1}$, to account for the parts that intersect the jth and $(j+1)$th domain walls. Given J', the choice of J for this decomposition is unique but there may be multiple choices of (K, K') for which it holds. Note also that a trajectory in

$\mathcal{B}_{H,j}$ can be decomposed into individual domain-wall pairs that coincide with the jth domain wall of H at some time step t and annihilate at some time step t'. The combined weight of all such pairs, given fixed coincidence point at $e_{H,j,t}$, is at most $(2q/(q^2+1))^{2t'}$. If we sum over $t' \geq t$, we find that the combined weight for all possible domain-wall pairs coinciding at time step t is at most

$$\frac{\left(\frac{2q}{q^2+1}\right)^{2t}}{1-\left(\frac{2q}{q^2+1}\right)^2} = \frac{(q^2+1)^2}{(q^2-1)^2}\left(\frac{2q}{q^2+1}\right)^{2t}. \quad (D36)$$

There can be many domain-wall pairs that intersect the jth domain wall of H but for each value of t, there will either be no intersection (in which case the factor is 1) or one intersection (in which case the factor is at most the above quantity). Thus we can take the product over including or not including a domain wall at each value of t and find

$$\sum_{K \in \mathcal{B}_{H,j}} \text{weight}(K) \leq \prod_{t=1}^{s}\left[1+\frac{(q^2+1)^2}{(q^2-1)^2}\left(\frac{2q}{q^2+1}\right)^{2t}\right]. \quad (D37)$$
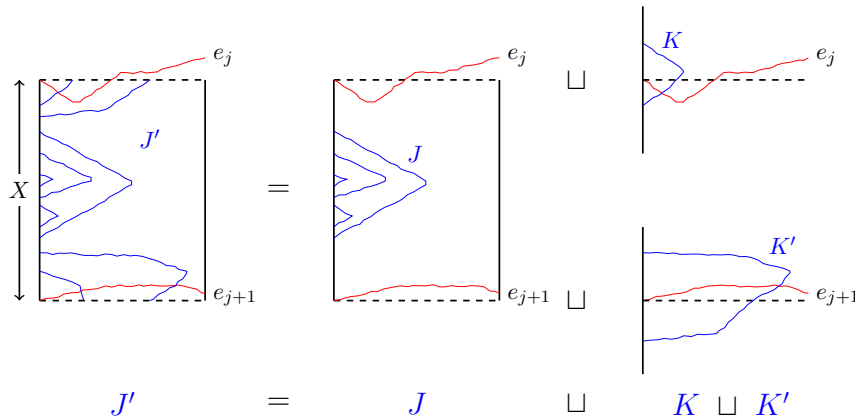
This implies that

FIG. 7. The outline of the argument in the proof of Lemma 4 that the sum over domain-wall trajectories in $\mathcal{J}_{H,j}$ is at least the sum in $\mathcal{J}'_{H,j,0}$ divided by some constant factor, expressed in Eq. (D38). Every trajectory in $J' \in \mathcal{J}'_{H,j,0}$ can be decomposed into a trajectory $J \in \mathcal{J}_{H,j}$ a trajectory $K \in \mathcal{B}_{H,j}$, and a trajectory $K' \in \mathcal{B}_{H,j+1}$, where each domain wall in K intersects the jth domain wall of H, and each domain wall of K' intersects the $(j+1)$th domain wall of H. Because the combined weight of all possible K and K' is only a constant factor, independent of n, the combined weight of all possible J cannot be more than a constant factor smaller than the combined weight of all possible J'. Note that the system with X sites has periodic boundary conditions in this figure.

$$\sum_{J' \in \mathcal{J}'_{H,j,0}} \text{weight}(J) \leq \left[\sum_{J \in \mathcal{J}_{H,j}} \text{weight}(J) \right] \left[\sum_{K \in \mathcal{B}_{H,j}} \text{weight}(K) \right] \left[\sum_{K' \in \mathcal{B}_{H,j+1}} \text{weight}(K') \right]$$

$$\leq \left[\sum_{J \in \mathcal{J}_{H,j}} \text{weight}(J) \right] \prod_{t=1}^{s} \left[1 + \frac{(q^2+1)^2}{(q^2-1)^2} \left(\frac{2q}{q^2+1} \right)^{2t} \right]^2$$

$$\leq \left[\sum_{J \in \mathcal{J}_{H,j}} \text{weight}(J) \right] \exp \left[2 \sum_{t=1}^{s} \frac{(q^2+1)^2}{(q^2-1)^2} \left(\frac{2q}{q^2+1} \right)^{2t} \right]$$

$$\leq \left[\sum_{J \in \mathcal{J}_{H,j}} \text{weight}(J) \right] \exp \left[\frac{8q^2}{(q^2-1)^2} \frac{1}{1 - \left(\frac{2q}{q^2+1} \right)^2} \right]$$

$$= \left[\sum_{J \in \mathcal{J}_{H,j}} \text{weight}(J) \right] \exp \left[\frac{8q^2(q^2+1)^2}{(q^2-1)^4} \right]$$

$$\leq \left[\sum_{J \in \mathcal{J}_{H,j}} \text{weight}(J) \right] e^{10}, \tag{D38}$$

where the last inequality follows since $q \geq 2$ and the function of q inside the exp is monotonically decreasing. By combining the above with Eq. (D35), we arrive at

$$\left(\sum_{J \in \mathcal{J}_{H,j}} \text{weight}(J) \right)$$
$$\geq \left(\frac{q+1}{q} \right)^X \left(\frac{1}{3e^{10}} \right) = \left(\frac{q+1}{q} \right)^X c^{-1}. \tag{D39}$$

∎

APPENDIX E: BOUNDS FOR THE COMPLETE-GRAPH ARCHITECTURE

1. Proof intuition and guide

In the following sections, we complete the proofs for upper and lower bounds of the complete-graph architecture, defined formally in Definition 1. The first insight about the complete-graph architecture is that all configurations with the same Hamming weight are equivalent, as there is a symmetry upon permutation of the qudits. Thus, trajectories through configuration space $\{I, S\}^n$ are reduced to trajectories through Hamming weight space $\{0, 1, \ldots, n\}$.

Our upper bound uses the framework of the unbiased walk and the lower bound uses the biased walk. Recall that we can use the unbiased walk to express the collision probability Z as a sum over all possible paths that the trajectory

might take, working from Eq. (B19):

$$Z = \frac{1}{(q+1)^n} \sum_{\vec{\gamma}^{(0)}} \mathop{\mathbb{E}}_{P_u, \vec{\gamma}^{(0)}} \left[\left(\frac{2q}{q^2+1} \right)^{\left(\begin{subarray}{c} \text{number of bit} \\ \text{flips during walk} \end{subarray} \right)} \right]$$
$$= \frac{1}{(q+1)^n} \sum_{x=0}^{n} \binom{n}{x} \mathop{\mathbb{E}}_{P_{u,x}} \left[\left(\frac{2q}{q^2+1} \right)^{\left(\begin{subarray}{c} \text{number of bit} \\ \text{flips during walk} \end{subarray} \right)} \right], \tag{E1}$$

where the $\binom{n}{x}$ comes from the fact that this is the number of initial configurations with Hamming weight x. For the complete-graph case, P_u takes on a simple form: if the current configuration is x, the chance that the configuration changes on the next time step is precisely the chance of finding mismatching values upon drawing a random pair of indices in $[n]$, which is given by $2x(n-x)/n(n-1)$, and if it does change, it is equally likely to become $x-1$ or to become $x+1$. For the biased walk P_b, everything is the same except that when the configuration changes, it is biased to travel to $x-1$ with probability $q^2/(q^2+1)$. Also, in the biased case, the initial configuration is not a uniform choice over all configurations but instead distributed according to Λ_b and the expectation in the above equation is replaced with $\mathbb{E}[q^{|\vec{\gamma}^{(s)}|}]$. Thus larger Hamming weight configurations are exponentially more significant in their contribution to Z.

To gain an intuition for what we expect, we first think about the biased walk, which is what we use for the lower bound. Here, the peak of the probability mass in the initial configuration Λ_b starts around Hamming weight $x = n/(q+1)$. On average, the walk lingers for $n(n-1)/2x(n-x)$ time steps before moving, which is approximately equal to $n/2x$ when x is close to 0. Due to the bias, and due to the time required to wait, the effective speed of the biased walk is

$$v_{\text{eff}}(x) = \frac{2x(n-x)}{n(n-1)} \left(\frac{q^2}{q^2+1} - \frac{1}{q^2+1} \right) \approx \frac{2x}{n} \frac{q^2-1}{q^2+1}$$
(E2)

in the direction of 0, since each time it moves, it has a $q^2/(q^2+1)$ chance of moving one unit closer to 0 but a $1/(q^2+1)$ chance of moving one unit further away from 0. Thus, in expectation, the time it takes for the peak of the probability mass to reach value 0 is

$$\sum_{x=1}^{n/(q+1)} \frac{1}{v_{\text{eff}}(x)} \approx \frac{q^2+1}{q^2-1} \frac{1}{2} n \log(n) \approx: s^*,$$
(E3)

noting that $\sum_x (1/x) \approx \log(n)$.

This strongly suggests that s^* time steps are *necessary* for anticoncentration, as any less time would mean that the peak of the distribution over Hamming weights at the end of the circuit would be located at some Hamming weight $y > 0$ and as a result, it would receive a significant amount of weight q^y in its contribution to Z. This is the intuition for our lower bound.

The biased walk also gives intuition for why there is a matching upper bound. If the circuit size is a little bigger than the lower bound, we expect the peak of the distribution to have terminated at the fixed point at 0. It is still possible that the tail of the distribution, which will not yet have reached the fixed point, is too fat for anticoncentration to have been achieved; each unit further away from 0 results in a factor-of-q-larger contribution to Z, so we need the tail to be exponentially decaying if we want to be able to ignore it. This is essentially what we are able to show, albeit in such a way that it might not be completely clear that this is what we are doing. Intuitively, one reason we expect this exponentially decaying tail is because the effective speed slows down as we get closer to zero. This gives the tail of the distribution, which is sitting further away from 0, time to "catch up," as its effective speed is faster.

To actually perform the upper bound, we turn back to the unbiased walk. To be clear and to match the progression in the full proof, we introduce the concept of a *reduced* path (equivalently, "reduced walk") as a walk that never stands still at a certain configuration. If its Hamming weight at time step t is y, then its Hamming weight at time step $t+1$

will move to $y-1$ or $y+1$. For each walk, we can form a corresponding reduced walk simply by removing consecutive duplicates from the sequence of configurations. Another way to look at it is that given a fixed reduced walk, the actual walk will linger at each location for a certain number of time steps before continuing. In the limit of large circuit size s, there is enough time for the actual walk to linger as long as it would like at each step and any reduced walk will successfully be "completed" by the actual walk. In this limit, $Z = Z_H$ where $Z_H := 2/(q^n+1)$ is the Haar value. Away from this limit, there is some probability that some reduced paths will not be completed.

We are able to express the difference between Z and Z_H as a sum over all reduced paths, including reduced paths that do not terminate at Hamming weight 0 or Hamming weight n, where the summand is proportional to the probability that the reduced path is not completed within s time steps:

$$(q+1)^n (Z - Z_H)$$
$$= \sum_{\text{red path } \phi} \left(\frac{q}{q^2+1} \right)^{\text{length of } \phi} \frac{(q-1)^2}{2q}$$
$$\times \Pr[\phi \text{ not completed in } s \text{ time steps}].$$
(E4)

The next key insight is to use a Chernoff bound to bound the probability of a reduced walk not being completed. If L is the length of the walk, the Chernoff bound states (for any constant $a > 0$) that

$$\Pr[L > s] \leq \frac{\mathbb{E}[e^{aL}]}{e^{as}},$$
(E5)

but this is particularly useful because, for fixed ϕ, L is itself a sum of independent random variables $L_{\phi^{(i)}}$, the number of time steps the walk waits on step i. Thus

$$\mathbb{E}[e^{aL}] = \prod_i \mathbb{E}[e^{aL_{\phi^{(i)}}}]$$
(E6)

and because each random variable $L_{\phi^{(i)}}$ is exponentially distributed, we can calculate $\mathbb{E}[e^{aL_{\phi^{(i)}}}]$ exactly. For the purposes of the proof sketch, denote

$$T_x := \mathbb{E}[e^{aL_x}].$$
(E7)

The dependence of T_x on x (the walk will wait longer when x is near 0 or n than when it is near $n/2$) appears to be a problem, as it is unclear how to actually perform the following sum over all possible ϕ (where the constant a for

each reduced walk ϕ, denoted a_ϕ, is specified later):

$$(q+1)^n(Z - Z_H)$$

$$= \frac{(q-1)^2}{2q} \sum_{\text{red path } \phi} e^{-a_\phi s} \prod_{i=1}^{\text{length of } \phi} \left(\frac{q}{q^2+1} T_{\phi^{(i)}} \right). \tag{E8}$$

To proceed, we break ϕ up into subpaths that inch closer and closer to 0 and n. We can write ϕ as the concatenation of $\phi_x, \phi_{x-1}, \ldots, \phi_w$, where ϕ_v begins at either v or $n - v$ and only reaches $v - 1$ or $n - v + 1$ for the first time on the very last step. Then, w is the minimum Hamming weight distance from one of the fixed points (0 or n) that the reduced walk ϕ ever reaches. Because the walk ϕ_v spends all its time between v and $n - v$, the expectation T_y for all of the y within one of these ϕ_v walks will be less than or equal to T_v (the walk moves slower when it is closer to 0 or n) and we can write

$$(q+1)^n(Z - Z_H) \le \frac{(q-1)^2}{2q} \sum_{x=0}^{n} \binom{n}{x} \sum_{w=0}^{\min(x, n-x)} e^{-a_w s}$$

$$\times \left[\sum_{\phi_x} \left(\frac{q}{q^2+1} T_x \right)^{\text{length of } \phi_x} \right] \cdots$$

$$\times \left[\sum_{\phi_w} \left(\frac{q}{q^2+1} T_w \right)^{\text{length of } \phi_w} \right]. \tag{E9}$$

Here, the $\binom{n}{x}$ comes in as the number of configurations with Hamming weight x and a_ϕ has changed to a_w because we choose it so that it only depends on the end point w of ϕ.

The above equation is huge progress, because we already know how to perform the sums in brackets. Essentially, the factor of T_x simply changes the *effective* value of q; we may define $\bar{q}$ to satisfy

$$\bar{q}/(\bar{q}^2 + 1) = q T_x/(q^2 + 1). \tag{E10}$$

Then, we can use the formulas for sums over paths that we have already developed in Lemma 1 to perform the sums. What we find is that, for the values of a_w that we can choose, we must allocate roughly $(q^2 + 1)n/2(q^2 - 1)x$ time steps for s such that $e^{-a_w s}$ can cancel out the value of the sum in brackets for ϕ_x. Note that this is precisely the inverse of the effective speed that we defined before. Then, for all the sums to be canceled from $v = 1$ to $v = n/2$, we must allocate

$$\sum_{v=1}^{n/2} \frac{q^2+1}{2(q^2-1)} \frac{n}{x} \approx \frac{q^2+1}{2(q^2-1)} n \log(n) \tag{E11}$$

time steps. Fundamentally, the $\log(n)$ factor becomes necessary because the walk waits longer and longer as it gets closer and closer to the fixed points. In the full analysis, we find that a term linear in n is also necessary to fully anticoncentrate, but our analysis of the linear term is not tight.

2. Preliminaries

a. Trajectories

For the complete-graph architecture, we may keep track of only the Hamming weight of a certain configuration. Thus, our random walks are over the set $\{0, 1, \ldots, n\}$. A trajectory γ is now a sequence of integers $(\gamma^{(0)}, \ldots, \gamma^{(s)})$. Generally speaking, if $t > s$ for a sequence of length s, let $\gamma^{(t)}$ return $\gamma^{(s)}$. A sequence is valid if for every t, $|\gamma^{(t)} - \gamma^{(t-1)}| \le 1$ and such that if 0 or n appears, it appears only once at the very end of the sequence. Let Γ be the set of all valid trajectories.

For any valid trajectory γ, the unbiased random walk associates a nonzero probability:

$$\Pr_{P_u}[\gamma] := \prod_{t=1}^{s} P_u[\gamma^{(t)} | \gamma^{(t-1)}], \tag{E12}$$

where

$$P_u[y|x] = \begin{cases} \dfrac{x(n-x)}{n(n-1)}, & \text{if } |y - x| = 1, \\ 1 - \dfrac{2x(n-x)}{n(n-1)}, & \text{if } y = x, \\ 0, & \text{otherwise.} \end{cases} \tag{E13}$$

We can make the same definition for the biased random walk by replacing P_u with P_b, where

$$P_b[y|x] = \begin{cases} \dfrac{2q^{1+x-y}}{q^2+1} \dfrac{x(n-x)}{n(n-1)}, & \text{if } |y - x| = 1, \\ 1 - \dfrac{2x(n-x)}{n(n-1)}, & \text{if } y = x, \\ 0, & \text{otherwise.} \end{cases} \tag{E14}$$

For $P \in \{P_u, P_b\}$ and any subset $\Upsilon \subseteq \Gamma$, we let $\Pr_P[\Upsilon] = \sum_{\gamma \in \Upsilon} \Pr_P[\gamma]$ be the total probability assigned to paths in Υ.

b. Conditional probabilities and expectations

For any $\gamma \in \Upsilon$, we may also define the conditional probability

$$\Pr_P[\gamma | \Upsilon] := \frac{\Pr_P[\gamma]}{\Pr_P[\Upsilon]}, \tag{E15}$$

which indicates drawing from the subset Υ with a probability proportional to that assigned by the (unbiased or

biased) random walk. This also allows us to naturally define conditional expectation values for some quantity Q computed from γ:

$$\mathbb{E}_P[Q[\gamma]|\gamma \in \Upsilon] := \sum_{\gamma \in \Upsilon} \frac{\Pr_P[\gamma]}{\Pr_P[\Upsilon]} Q[\gamma]. \qquad (E16)$$

c. Trajectory concatenation and other operations

For any trajectory $\gamma = (\gamma^{(0)}, \ldots, \gamma^{(s)}) \in \Gamma$, let $L[\gamma] = s$ be the length of the trajectory and let $\gamma^{(L)}$ be shorthand for $\gamma^{(L[\gamma])}$. The statement $w \in \gamma$ returns true if there exists some t for which $\gamma^{(t)} = w$. Then, we let

$$S_w[\gamma] = \begin{cases} \min \left(\{t : \gamma^{(t)} = w\}\right), & \text{if } w \in \gamma, \\ -1, & \text{if } w \notin \gamma, \end{cases} \qquad (E17)$$

be the first time step along γ for which the trajectory reaches w. We also let

$$M[\gamma] := \max_{0 \leq t \leq L[\gamma]} \gamma_t, \qquad (E18)$$

$$m[\gamma] := \min_{0 \leq t \leq L[\gamma]} \gamma_t \qquad (E19)$$

be the maximum and minimum Hamming weights through which the trajectory passes.

We can naturally concatenate two trajectories γ_1 and γ_2 if $\gamma_1^{(L)} = \gamma_2^{(0)}$ to form a trajectory $\gamma = \gamma_1 \cdot \gamma_2$ of length $L[\gamma_1] + L[\gamma_2]$. We say that $\gamma_A \subset \gamma_C$ if there exists some γ_B for which $\gamma_A \cdot \gamma_B = \gamma_C$.

For any trajectory γ, we let $\tilde{\gamma}$ be the flipped trajectory:

$$\tilde{\gamma} := (n - \gamma^{(0)}, n - \gamma^{(1)}, \ldots, n - \gamma^{(s)}). \qquad (E20)$$

In general, if v is an integer with $0 \leq v \leq n$, then let $\tilde{v} := \min(v, n - v)$.

Similarly, let $\bar{\gamma}$ return the reversed trajectory:

$$\bar{\gamma} := (\gamma^{(s)}, \ldots, \gamma^{(0)}). \qquad (E21)$$

Moreover, let $\gamma^{[t]}$ return the trajectory γ truncated to length t, or simply return γ if $t \geq L[\gamma]$, i.e.,

$$\gamma^{[t]} := \begin{cases} \gamma, & \text{if } t \geq L[\gamma], \\ (\gamma^{(0)}, \ldots, \gamma^{(t)}), & \text{if } t < L[\gamma]. \end{cases} \qquad (E22)$$

Let $\gamma^{[L-t]}$ be shorthand for $\gamma^{[L[\gamma]-t]}$. More generally, let $\gamma^{[a,b]} = (\gamma^{(a)}, \gamma^{(a+1)}, \ldots, \gamma^{(b)})$.

d. Important subsets of Γ

We now define various subsets of Γ. Let

$$\Gamma_x := \{\gamma \in \Gamma : \gamma^{(0)} = x\} \qquad (E23)$$

be the subset of trajectories that begin at x and let

$$\Gamma^w := \{\gamma \in \Gamma : S_w[\gamma] = L[\gamma]\} \qquad (E24)$$

be the set of trajectories that reach w for the first time and immediately terminate. We make the natural combination of these:

$$\Gamma_x^w := \{\gamma \in \Gamma : \gamma_0 = x, S_w[\gamma] = L[\gamma]\}. \qquad (E25)$$

Of particular importance are sets where $w = 0$ or $w = n$, which include valid trajectories that terminate at one of the fixed points of the random walk. Define

$$\Gamma_x^* = \Gamma_x^0 \cup \Gamma_x^n \qquad (E26)$$

and note that $\Pr_{P_u}[\Gamma_x^*] = \Pr_{P_b}[\Gamma_x^*] = 1$, a statement that intuitively makes sense since walks will eventually reach either 0 or n with probability 1. Adding the superscript w to any set Υ restricts to walks for which $L[\gamma] = S_w[\gamma]$.

When any walk in Υ can be concatenated with any walk in Υ', we let

$$\Upsilon \cdot \Upsilon' := \{\gamma \cdot \gamma' : \gamma \in \Upsilon, \gamma' \in \Upsilon'\}. \qquad (E27)$$

Additionally, we let

$$\tilde{\Upsilon} := \{\tilde{\gamma} : \gamma \in \Upsilon\}, \qquad (E28)$$

$$\bar{\Upsilon} := \{\bar{\gamma} : \gamma \in \Upsilon\}. \qquad (E29)$$

e. Reduced trajectories

We also introduce the concept of a *reduced* trajectory, which we sometimes refer to synonymously as a reduced walk, which is a valid trajectory for which $|\gamma^{(t)} - \gamma^{(t-1)}| = 1$ for all t; that is, the reduced walk never stands still. We let the set of reduced walks be Ψ and we let all subscripts and superscripts restrict Ψ in the same way that they restrict Γ. For any $\gamma \in \Gamma$, we can associate a reduced walk $\psi \in \Psi$ by removing consecutive duplicates from γ. Under this definition, we let $R[\gamma] := \psi$. For any $\psi \in \Psi$, we let

$$\Gamma_\psi = \{\gamma \in \Gamma : R[\gamma] = \psi, R[\psi^{[L-1]}] \neq \psi\}, \qquad (E30)$$

where the second condition acts to include only trajectories γ the final configuration of which appears only once (i.e., when the final configuration is removed, the reduced sequence changes).

Under dynamics by either the unbiased or biased walk, it is easy to calculate the probability associated with Γ_ψ:

$$\Pr_{P_u}[\Gamma_\psi] := \left(\frac{1}{2}\right)^{L[\psi]},\tag{E31}$$

$$\Pr_{P_b}[\Gamma_\psi] := q^{\psi^{(0)}-\psi^{(L)}}\left(\frac{q}{q^2+1}\right)^{L[\psi]}.\tag{E32}$$

Finally, define the following subsets of Ψ:

$$\Lambda_x = \Psi^{x-1}_{x|n-x+1} \cup \Psi^{n-x+1}_{x|x-1},\tag{E33}$$

$$\Xi_w = \{\psi \in \Psi_w : m(\psi) \geq w, M(\psi) \leq n-w\},\tag{E34}$$

where the subset $\Psi^c_{a|b}$ is defined as follows:

$$\Psi^w_{x|z} := \begin{cases} \{\psi \in \Psi^w_x : M[\psi] < z\}, & \text{if } w < x < z, \\ \{\psi \in \Psi^w_x : m[\psi] > z\}, & \text{if } z < x < w, \\ \varnothing, & \text{otherwise.} \end{cases}\tag{E35}$$

In words, the set $\Psi^w_{x|z}$ includes reduced walks that begin at x and end at w without ever reaching z. Thus Λ_x is the set of reduced walks that start at x and end at $x-1$ without ever reaching $n-x+1$ or end at $n-x+1$ without ever reaching $x-1$. The set Ξ_w is the set of reduced walks of any finite length that start at w but never reach either $w-1$ or $n-w+1$.

3. Upper-bound proof

Theorem 11 (Theorem 2 from main text). *For the complete-graph architecture with circuit size s on n qudits with local Hilbert space dimension q,*

$$Z \leq Z_H\left[1 + e^{-\frac{2a}{n}(s-s^*)}\right],\tag{E36}$$

as long as $s \geq s^*$, where

$$s^* = \frac{q^2+1}{2(q^2-1)}n\log(n) + O(n),\tag{E37}$$

$$a = \frac{(q-1)^2}{2(q^2+1)}.\tag{E38}$$

Proof. In this proof, we are working with expressions for the collision probability Z. It will take several steps to manipulate the original expression into the form we need, so we will move back and forth between updating the expression and developing the tools needed to justify these updates.

We start by expressing

$$Z = \frac{1}{(q+1)^n}\sum_{x=0}^{n}\binom{n}{x}\mathop{\mathbb{E}}_{P_u}\left[\left(\frac{2q}{q^2+1}\right)^{L\left[R\left[\gamma^{[s]}\right]\right]} \middle| \gamma \in \Gamma^*_x\right].\tag{E39}$$

This is seen to be equivalent to Eq. (B19) as follows. There are $\binom{n}{x}$ initial configurations with Hamming weight x and generating a length-s trajectory beginning at x with the unbiased Markov chain is equivalent to randomly choosing a trajectory γ from Γ^*_x, which begins at x and ends at a fixed point (0 or n), with a probability proportional to that assigned by the unbiased walk, and then truncating the walk to length s, denoted by $\gamma^{[s]}$. Then, $R[\gamma^{[s]}]$ is the reduced trajectory, where consecutive duplicates are removed, and $L[R[\gamma^{[s]}]]$ is the length of that reduced trajectory or, in other words, the total number of bit flips that have occurred within the first s time steps.

Moving ahead, we observe that drawing γ from Γ^*_x is equivalent to first drawing a reduced trajectory ψ from Ψ^*_x and then drawing γ from Γ_ψ, so we can rewrite

$$Z = \frac{1}{(q+1)^n}\sum_{x=0}^{n}\binom{n}{x}\sum_{\psi \in \Psi^*_x}\Pr_{P_u}[\Gamma_\psi]\mathop{\mathbb{E}}_{P_u}$$
$$\times \left[\left(\frac{2q}{q^2+1}\right)^{L\left[R\left[\gamma^{[s]}\right]\right]} \middle| \gamma \in \Gamma_\psi\right].\tag{E40}$$

Now, note the following general statement about any integer-valued random variable X such that $0 \leq X \leq M$. For any function f, we have

$$\mathbb{E}[f(X)] = \sum_{m=0}^{M}\Pr[X = m]f(m)$$
$$= \sum_{m=0}^{M}(\Pr[X < m+1] - \Pr[X < m])f(m)$$
$$= f(M) + \sum_{m=1}^{M}\Pr[X < m]\left[f(m-1) - f(m)\right].\tag{E41}$$

Taking $X = L[R[\gamma^{[s]}]]$ for γ drawn at random from Γ_ψ and $f(X) = [2q/(q^2+1)]^X$, we find that

$$\underset{P_u}{\mathbb{E}}\left[\left(\frac{2q}{q^2+1}\right)^{L\left[R\left[\gamma^{[s]}\right]\right]}\middle|\gamma\in\Gamma_\psi\right]$$

$$=\left(\frac{2q}{q^2+1}\right)^{L[\psi]}+\sum_{m=1}^{L[\psi]}\underset{P_u}{\text{Pr}}\left[L[R[\gamma^{[s]}]]<m\middle|\gamma\in\Gamma_\psi\right]\left(\frac{2q}{q^2+1}\right)^{m-1}\frac{(q-1)^2}{q^2+1}$$

$$=\left(\frac{2q}{q^2+1}\right)^{L[\psi]}+\sum_{m=1}^{L[\psi]}\underset{P_u}{\text{Pr}}\left[L[\gamma]>s\middle|\gamma\in\Gamma_{\psi^{[m]}}\right]\left(\frac{2q}{q^2+1}\right)^{m-1}\frac{(q-1)^2}{q^2+1},\tag{E42}$$

where the last line follows since the conditions $L[R[\gamma^{[s]}]]<m$ with $\gamma\in\Gamma_\psi$ and $L[\gamma]>s$ with $\gamma\in\Gamma_{\psi^{[m]}}$ both correspond to deciding if the configuration has changed at least m times within the first s steps.

The quantity

$$\frac{1}{(q+1)^n}\sum_{x=0}^n\binom{n}{x}\sum_{\psi\in\Psi_x^*}\underset{P_u}{\text{Pr}}[\Gamma_\psi]\left(\frac{2q}{q^2+1}\right)^{L[\psi]}\tag{E43}$$

is precisely equal to Z_H, as this represents the limit of infinite size where all trajectories terminate at one of the fixed points (see Sec. 6). Thus, also noting that $\text{Pr}_{P_u}[\Gamma_\psi]=2^{-L[\psi]}$, we have

$$Z=Z_H+\frac{(q-1)^2}{(q+1)^n(q^2+1)}\sum_{x=0}^n\binom{n}{x}\sum_{\psi\in\Psi_x^*}\underset{P_u}{\text{Pr}}[\Gamma_\psi]\sum_{m=1}^{L[\psi]}\underset{P_u}{\text{Pr}}\left[L[\gamma]>s\middle|\gamma\in\Gamma_{\psi^{[m]}}\right]\left(\frac{2q}{q^2+1}\right)^{m-1}$$

$$=Z_H+\frac{(q-1)^2}{(q+1)^n(2q)}\sum_{x=0}^n\binom{n}{x}\sum_{\psi\in\Psi_x^*}\underset{P_u}{\text{Pr}}[\Gamma_\psi]\sum_{m=1}^{L[\psi]}\underset{P_u}{\text{Pr}}\left[L[\gamma]>s\middle|\gamma\in\Gamma_{\psi^{[m]}}\right]\left(\frac{2q}{q^2+1}\right)^{m}$$

$$=Z_H+\frac{(q-1)^2}{(q+1)^n(2q)}\sum_{x=0}^n\binom{n}{x}\sum_{\psi\in\Psi_x^*}2^{-L[\psi]}\sum_{m=1}^{L[\psi]}\underset{P_u}{\text{Pr}}\left[L[\gamma]>s\middle|\gamma\in\Gamma_{\psi^{[m]}}\right]\left(\frac{2q}{q^2+1}\right)^{m}.\tag{E44}$$

Now, in the first line below, by associating $\phi=\psi^{[m]}$ we reorder and regroup the sums: instead of summing over paths ψ that end at a fixed point and then all intermediate points $m=1,\ldots,L[\psi]$ along the path, we first sum over all m, all ϕ (not necessarily ending at a fixed point) of length m, and then all ψ for which $\phi\subset\psi$ (recall that this means that the first $L[\phi]$ entries in the trajectory ψ are equal to ϕ). In the second line, we note that the sums over m and ϕ of length m can be replaced by a simple sum over all ϕ (of any length). In the third line, we note that the total probability of all the walks ψ for which $\phi\subset\psi$ is just $2^{-L[\phi]}$:

$$Z=Z_H+\frac{(q-1)^2}{(q+1)^n(2q)}\sum_{x=0}^n\binom{n}{x}\sum_{\substack{m=1\\ }}^{\infty}\sum_{\substack{\phi\in\Psi_x\\ L[\phi]=m}}\left(\sum_{\substack{\psi\in\Psi_x^*\\ \phi\subset\psi}}2^{-L[\psi]}\right)\underset{P_u}{\text{Pr}}\left[L[\gamma]>s\middle|\gamma\in\Gamma_\phi\right]\left(\frac{2q}{q^2+1}\right)^{L[\phi]}$$

$$=Z_H+\frac{(q-1)^2}{(q+1)^n(2q)}\sum_{x=0}^n\binom{n}{x}\sum_{\phi\in\Psi_x}\left(\sum_{\substack{\psi\in\Psi_x^*\\ \phi\subset\psi}}2^{-L[\psi]}\right)\underset{P_u}{\text{Pr}}\left[L[\gamma]>s\middle|\gamma\in\Gamma_\phi\right]\left(\frac{2q}{q^2+1}\right)^{L[\phi]}$$

$$=Z_H+\frac{(q-1)^2}{(q+1)^n(2q)}\sum_{x=0}^n\binom{n}{x}\sum_{\phi\in\Psi_x}\left(2^{-L[\phi]}\right)\underset{P_u}{\text{Pr}}\left[L[\gamma]>s\middle|\gamma\in\Gamma_\phi\right]\left(\frac{2q}{q^2+1}\right)^{L[\phi]}$$

$$=Z_H+\frac{(q-1)^2}{(q+1)^n(2q)}\sum_{x=0}^n\binom{n}{x}\sum_{\phi\in\Psi_x}\underset{P_u}{\text{Pr}}\left[L[\gamma]>s\middle|\gamma\in\Gamma_\phi\right]\left(\frac{q}{q^2+1}\right)^{L[\phi]}.\tag{E45}$$

Now, we examine the final expression. The difference between Z and Z_H is a sum over Ψ_x, which includes all reduced paths ϕ that start at x and may or may not terminate at 0 or n. The statement $L[\gamma] > s$ is true if the number of time steps it takes to complete this reduced path is at least s, i.e., the probability that *the path does not finish within s time steps*. As a sanity check, when s becomes infinite, we expect this probability to become zero for any path, as there would be enough time for any path to finish, and in this case $Z = Z_H$ as expected. This expression represents progress because we will be able to bound the probability of a certain path being completed using a Chernoff bound.

For any random variable X and for any constant $a > 0$,

$$\Pr[X > k] \leq \frac{\mathbb{E}[e^{aX}]}{e^{ak}}. \tag{E46}$$

We use this bound with $X = L[\gamma]$, $k = s$ and yet-to-be-specified constants $a_\phi > 0$:

$$Z - Z_H \leq \frac{(q-1)^2}{(q+1)^n(2q)} \sum_{x=0}^{n} \binom{n}{x} \sum_{\phi \in \Psi_x} e^{-a_\phi s} \mathbb{E}_{P_u}$$

$$\times \left[e^{a_\phi L[\gamma]} \mid \gamma \in \Gamma_\phi \right] \left(\frac{q}{q^2+1} \right)^{L[\phi]}. \tag{E47}$$

The Chernoff bound has the additional benefit that $\mathbb{E}[e^{aX}]$ separates when X is the sum of independent random variables. In particular, once ϕ is fixed, $L[\gamma]$ is the sum of exponentially distributed random variables corresponding to how many time steps the path γ waits at each position along the reduced path ϕ.

This is seen formally by noting that

$$\phi = (\phi^{(0)}, \phi^{(1)}) \cdot (\phi^{(1)}, \phi^{(2)}) \cdot \ldots \cdot (\phi^{(L-1)}, \phi^{(L)}), \tag{E48}$$

$$\Gamma_\phi = \Gamma_{(\phi^{(0)},\phi^{(1)})} \cdot \Gamma_{(\phi^{(1)},\phi^{(2)})} \cdot \ldots \cdot \Gamma_{(\phi^{(L-1)},\phi^{(L)})}, \tag{E49}$$

and meanwhile, for any collection of subsets Υ_m,

$$\mathbb{E}_{P_u}\left[e^{aL[\gamma]} \mid \gamma \in \Upsilon_1 \cdot \ldots \cdot \Upsilon_M \right] = \prod_{m=1}^{M} \mathbb{E}_{P_u}\left[e^{aL[\gamma]} \mid \gamma \in \Upsilon_m \right]. \tag{E50}$$

For any $r = 0, \ldots, L[\phi] - 1$, we can evaluate

$$\mathbb{E}_{P_u}\left[e^{aL[\gamma]} \mid \gamma \in \Gamma_{(\phi^{(r)},\phi^{(r+1)})} \right] = \sum_{t=1}^{\infty} \left(1 - \lambda_{\phi^{(r)}}^{-1} \right)^{t-1} \lambda_{\phi^{(r)}}^{-1} e^{at}$$

$$= \frac{1}{1 - \lambda_{\phi^{(r)}}(1 - e^{-a})}, \tag{E51}$$

where

$$\lambda_v := \frac{n(n-1)}{2v(n-v)} \tag{E52}$$

is the expected amount of time the walk will wait at Hamming weight v before moving to $v+1$ or $v-1$, and hence

$$\mathbb{E}_{P_u}\left[e^{a_\phi L[\gamma]} \mid \gamma \in \Gamma_\phi \right] = \prod_{r=0}^{L[\phi]-1} \frac{1}{1 - \lambda_{\phi^{(r)}}(1 - e^{-a_\phi})}. \tag{E53}$$

We have made some progress at evaluating the bound on Z but at this point it remains unclear how to perform the sum over paths $\phi \in \Psi_x$. To do so, first we decompose paths ϕ into a series of subpaths that inch closer and closer to the fixed points at 0 and n. In particular, we decompose a path ϕ as a concatenation of subpaths drawn from Λ_v for various v and one final subpath drawn from Ξ_w, as described in the following lemma. Recall from Eqs. (E33) and (E34) that these subsets of Ψ are defined by where they start, where they end, and the maximum or minimum point they reach.

Lemma 5. *Suppose that $\phi \in \Psi_x$. Let $\tilde{x} = \min(x, n-x)$ and let $w = \min[m(\phi), n - M(\phi)]$. Then there is a unique sequence of trajectories $(\phi_v)_{v=w}^{\tilde{x}}$ with $\phi_v \in \Lambda_v$ for $v = w+1, \ldots, \tilde{x}$ and $\phi_w \in \Xi_w$ and such that*

$$\phi = \alpha_{\tilde{x}} \cdot \alpha_{\tilde{x}-1} \cdot \ldots \cdot \alpha_w, \tag{E54}$$

where for each v, either $\alpha_v = \phi_v$ or $\alpha_v = \tilde{\phi}_v$, depending on whether α_{v+1} terminates at v or at $n - v$.

Proof. Let r_v be the minimum r such that $\phi^{(r)} = v$ or $\phi^{(r)} = n - v$. Then, for each $v = w+1, \ldots, \tilde{x}$, we can define

$$\alpha_v = \phi^{[r_v, r_{v-1}]} \tag{E55}$$

and

$$\alpha_w = \phi^{[r_w, L[\phi]]}. \tag{E56}$$

Then, each α_v begins at either v or $n - v$ and terminates upon reaching either $v - 1$ or $n - v + 1$ for the first time. Hence it is a member of Λ_v or $\tilde{\Lambda}_v$ but not both. Finally, α_w is a member of Ξ_w because it begins at w and never reaches either $w - 1$ or $n - w + 1$, since this would contradict the definition of w. ∎

We use the notation $\tilde{v} := \min(v, n - v)$ for any integer v throughout the remainder of the proof. The above lemma allows us to replace the sum over $\phi \in \Psi_x$ with sums over w from 0 to $\tilde{x}$ and sums over $\phi_v \in \Lambda_v$, $\phi_w \in \Xi_w$. The summand is a product of factors $[1 - \lambda_{\phi^{(r)}}(1 - e^{-a_\phi})]^{-1}$, each of which can be collected within just one of the sums. Moreover, the fact that these products are invariant under reversing the path, i.e.,

$$\prod_{r=0}^{L[\psi]-1} f\,[\psi^{(r)}] = \prod_{r=0}^{L[\psi]-1} f\,[\tilde{\psi}^{(r)}], \tag{E57}$$

for any function f, means that it is unimportant that α_v can equal ϕ_v or $\tilde{\phi}_v$ as both yield the same result.

We choose a_ϕ so that it only depends on $w = \min[m(\phi), n - M(\phi)]$, denoted henceforth by a_w. Collecting these observations and noting that the $L[\phi]$ factors of $q/(q^2 + 1)$ can each be allocated to one of the steps taken in ϕ, we find that

$$\frac{(q+1)^n(2q)}{(q-1)^2}\,(Z - Z_H) \leq \sum_{x=0}^{n} \binom{n}{x} \sum_{w=0}^{\tilde{x}} e^{-a_w s} \prod_{v=w}^{\tilde{x}} \left[\sum_{\substack{\phi_v \in \Lambda_v \\ \text{or} \\ \phi_w \in \Xi_w}} \prod_{r=0}^{L[\phi_v]-1} \frac{q}{q^2+1} \frac{1}{1 - \lambda_{\phi_v^{(r)}}(1 - e^{-a_w})} \right]$$

$$\leq \sum_{x=0}^{n} \binom{n}{x} \sum_{w=0}^{\tilde{x}} e^{-a_w s} \prod_{v=w}^{\tilde{x}} \left\{ \sum_{\substack{\phi_v \in \Lambda_v \\ \text{or} \\ \phi_w \in \Xi_w}} \left[\frac{q}{q^2+1} \frac{1}{1 - \lambda_v(1 - e^{-a_w})} \right]^{L[\phi_v]} \right\}, \tag{E58}$$

where in the final line, we use the fact that by definition of $\phi_v \in \Lambda_v$ or $\phi_v \in \Xi_v$, $v \leq \phi_v^{(r)} \leq n - v$ for all $r < L[\phi]$ and also $\lambda_v \geq \lambda_{v'}$ whenever $\tilde{v} < \tilde{v}'$.

This form is very useful because we know how to perform sums in parentheses, using the strategy we have first seen in Lemma 1. The values of these sums are given by the following lemma, the proof of which is delayed until after the main proof.

Lemma 6. *Given v and a parameter a that satisfies $1 \leq e^a \leq \{1 - [(q-1)^2/q^2 + 1](1/\lambda_v)\}^{-1}$,*

$$\sum_{\alpha \in \Lambda_v} \left[\frac{q}{q^2+1} \frac{1}{1 - \lambda_v(1 - e^{-a})} \right]^{L[\alpha]} = \bar{q}_{v,a}^{-1} \frac{1 + \bar{q}_{v,a}^{-n+2v}}{1 + \bar{q}_{v,a}^{-n+2v-2}}, \tag{E59}$$

$$\sum_{\alpha \in \Xi_v} \left[\frac{q}{q^2+1} \frac{1}{1 - \lambda_v(1 - e^{-a})} \right]^{L[\alpha]} = \frac{\bar{q}_{v,a}^2 + 1}{(\bar{q}_{v,a} - 1)^2} \left(1 - \bar{q}_{v,a}^{-1} \frac{1 + \bar{q}_{v,a}^{-n+2v}}{1 + \bar{q}_{v,a}^{-n+2v-2}} \right), \tag{E60}$$

where $\bar{q}_{v,a}$ is defined in Definition 6.

Definition 6. *Given x and parameter a satisfying $1 \leq e^a \leq (1 - ((q-1)^2/q^2 + 1)(1/\lambda_x))^{-1}$, let*

$$\bar{q}_{x,a} = \left(\frac{q^2+1}{2q} \right) [1 - \lambda_x(1 - e^{-a})] \left\{ 1 + \sqrt{1 - \frac{4q^2}{(q^2+1)^2[1 - \lambda_x(1 - e^{-a})]^2}} \right\}. \tag{E61}$$

Note that $\bar{q}_{x,a}$ satisfies the equation

$$\frac{\bar{q}_{x,a}}{\bar{q}_{x,a}^2 + 1} = \frac{q}{q^2+1} \frac{1}{1 - \lambda_x(1 - e^{-a})} \tag{E62}$$

and that

$$\bar{q}_{x,a}^{-1} = \left(\frac{q^2+1}{2q} \right) [1 - \lambda_x(1 - e^{-a})] \left\{ 1 - \sqrt{1 - \frac{4q^2}{(q^2+1)^2[1 - \lambda_x(1 - e^{-a})]^2}} \right\}. \tag{E63}$$

This lemma allows us to state

$$\frac{(q+1)^n(2q)}{(q-1)^2}(Z-Z_H) \leq \sum_{x=0}^{n}\binom{n}{x}\sum_{w=0}^{\tilde{x}}e^{-sa_w}\left(\prod_{v=w+1}^{\tilde{x}}\bar{q}_{v,a_w}^{-1}\frac{1+\bar{q}_{v,a_w}^{-n+2v}}{1+\bar{q}_{v,a_w}^{-n+2v-2}}\right)\left[\frac{\bar{q}_{w,a_w}^2+1}{(\bar{q}_{w,a_w}-1)^2}\left(1-\bar{q}_{w,a_w}^{-1}\frac{1+\bar{q}_{w,a_w}^{-n+2w}}{1+\bar{q}_{w,a_w}^{-n+2w-2}}\right)\right]$$

$$\leq \sum_{x=0}^{n}\binom{n}{x}\sum_{w=0}^{\tilde{x}}e^{-sa_w}\left(\prod_{v=w+1}^{\tilde{x}}\tilde{q}_{v,a_w}^{-1}\right)e^{\frac{q^2}{q^2-1}}\left[\frac{\bar{q}_{w,a_w}^2+1}{(\bar{q}_{w,a_w}-1)^2}\left(1-\bar{q}_{w,a_w}^{-1}\frac{1+\bar{q}_{w,a_w}^{-n+2w}}{1+\bar{q}_{w,a_w}^{-n+2w-2}}\right)\right], \quad \text{(E64)}$$

where the second line follows from the observation that (also noting $\bar{q}_{v,a} < q$ for all v, a)

$$\prod_{v=w+1}^{\tilde{x}}\frac{1+\bar{q}_{v,a_w}^{-n+2v}}{1+\bar{q}_{v,a_w}^{-n+2v-2}} \leq \prod_{m=0}^{\infty}\frac{1+q^{-2m}}{1+q^{-2m-2}} = \prod_{m=0}^{\infty}\left(1+q^{-2m}\frac{1-q^{-2}}{1+q^{-2m-2}}\right),$$

$$\leq \prod_{m=0}^{\infty}\left(1+q^{-2m}\right) \leq \prod_{m=0}^{\infty}\exp\left(q^{-2m}\right) = \exp\left(\sum_{m=0}^{\infty}q^{-2m}\right) = \exp\left(\frac{q^2}{q^2-1}\right). \quad \text{(E65)}$$

To continue, we make choices for a_w and show upper bounds for the various factors in the above expression. For $w > 0$, we make the specification for a_w that

$$\eta_w := 1 - e^{-a_w} := \frac{(q-1)^2}{q^2+1}\frac{1}{2\lambda_w} = \frac{(q-1)^2}{q^2+1}\frac{w(n-w)}{n(n-1)} \quad \text{(E66)}$$

and that $\eta_0 = \eta_1$. This choice implies that

$$a_w \geq \frac{(q-1)^2}{q^2+1}\frac{1}{2\lambda_w} = \frac{(q-1)^2}{q^2+1}\frac{w(n-w)}{n(n-1)}. \quad \text{(E67)}$$

Moreover, it implies that, so long as $w \leq x \leq n-w$,

$$1 - \lambda_x(1-e^{-a_w}) \geq 1 - \lambda_w(1-e^{-a_w}) = 1 - \frac{(q-1)^2}{2(q^2+1)} = \frac{(q+1)^2}{2(q^2+1)}, \quad \text{(E68)}$$

which, by Definition 6, implies that

$$q \geq \bar{q}_{x,a_w} \geq \frac{(q+1)^2}{4q}. \quad \text{(E69)}$$

When this is the case, we have

$$\left[\frac{\bar{q}_{w,a_w}^2+1}{(\bar{q}_{w,a_w}-1)^2}\left(1-\bar{q}_{w,a_w}^{-1}\frac{1+\bar{q}_{w,a_w}^{-n+2w}}{1+\bar{q}_{w,a_w}^{-n+2w-2}}\right)\right] \leq \frac{\bar{q}_{w,a_w}^2+1}{(\bar{q}_{w,a_w}-1)^2}\left(1-\bar{q}_{w,a_w}^{-1}\right)$$

$$= \frac{\bar{q}_{w,a_w}^2+1}{(\bar{q}_{w,a_w}-1)\bar{q}_{w,a_w}}$$

$$\leq \frac{q^2+1}{\left[\frac{(q+1)^2}{4q}-1\right]\frac{(q+1)^2}{4q}}$$

$$= \frac{(q^2+1)(4q)^2}{(q-1)^2(q+1)^2}$$

$$\leq \frac{320}{9}, \quad \text{(E70)}$$

where the last line follows for all $q \geq 2$, which will be true for any physically realizable circuit. This takes care of the final factor in Eq. (E64). What remains are the factors of $\bar{q}_{v,a_w}^{-1}$. To handle these, we use the following bound, the proof of which is delayed to the next section.

Lemma 7. *With $\bar{q}_x$ defined as in Definition 6 and as long as $1 \leq e^a \leq \{1 - [(q-1)^2/q^2 + 1](1/\lambda_x)\}^{-1}$,*

$$\bar{q}_{x,a}^{-1} \leq q^{-1} \exp\left\{a\left[\lambda_x \frac{q^2+1}{q^2-1} + \lambda_x^2(1 - e^{-a})\frac{(q^2+1)^4}{(q^2-1)^3}\right]\right\}. \tag{E71}$$

We also need the following observation, which holds under the assumption that $1 \leq j < k \leq (n/2)$:

$$\sum_{r=j+1}^{k} \lambda_r \leq \frac{n}{2}\sum_{r=j+1}^{k}\left(\frac{1}{r} + \frac{1}{n-r}\right) \leq \frac{n}{2}\left(\int_j^k d\rho\,\frac{1}{\rho} + \int_{n-k-1}^{n-j-1} d\rho\,\frac{1}{\rho}\right)$$

$$= \frac{n}{2}\left\{\log(k/j) + \log[(n-j-1)/(n-k-1)]\right\}$$

$$\leq \frac{n}{2}[\log(k/j) + \log(2)] < \frac{n}{2}[\log(2k/j) + 1], \tag{E72}$$

$$\sum_{r=j+1}^{k} \lambda_r^2 \leq \frac{n^2}{4}\sum_{r=j+1}^{k}\frac{n^2}{r^2(n-r)^2} \leq n^2\sum_{r=j+1}^{k}\frac{1}{r^2} \leq \frac{n^2}{j} < \frac{\pi^2 n^2}{6j}. \tag{E73}$$

Similarly, for the case where $j = 0$, we find that

$$\sum_{r=1}^{k} \lambda_r \leq \frac{n}{2}\big[\log(2k) + 1\big], \qquad \sum_{r=1}^{k} \lambda_r^2 \leq \frac{\pi^2 n^2}{6}. \tag{E74}$$

Now, we can write the following, where $\bar{w} := \max(w, 1)$:

$$\exp\left(-\frac{q^2}{q^2-1}\right)\frac{9(q+1)^n(2q)}{320(q-1)^2}(Z - Z_H)$$

$$\leq \sum_{x=0}^{n}\binom{n}{x}\sum_{w=0}^{\tilde{x}}e^{-sa_w}\left(\prod_{v=w+1}^{\tilde{x}}\bar{q}_{v,a_w}^{-1}\right)$$

$$\leq \sum_{x=0}^{n}\binom{n}{x}\sum_{w=0}^{\tilde{x}}e^{-sa_w}\left\{\prod_{v=w+1}^{\tilde{x}}q^{-1}\exp\left[\frac{q^2+1}{q^2-1}a_w\lambda_v + \frac{(q^2+1)^4}{(q^2-1)^3}a_w\eta_w\lambda_v^2\right]\right\}$$

$$= \sum_{x=0}^{n}\binom{n}{x}\sum_{w=0}^{\tilde{x}}e^{-sa_w}q^{-\tilde{x}+w}\exp\left[\frac{q^2+1}{q^2-1}a_w\sum_{v=w+1}^{\tilde{x}}\lambda_v + \frac{(q^2+1)^4}{(q^2-1)^3}a_w\eta_w\sum_{v=w+1}^{\tilde{x}}\lambda_v^2\right]$$

$$\leq \sum_{x=0}^{n}\binom{n}{x}\sum_{w=0}^{\tilde{x}}e^{-sa_w}q^{-\tilde{x}+w}\exp\left[\frac{q^2+1}{q^2-1}a_w\frac{n}{2}\left(\log\frac{2\tilde{x}}{\bar{w}}+1\right) + \frac{(q^2+1)^4}{(q^2-1)^3}a_w\eta_w\frac{n^2\pi^2}{6\bar{w}}\right]$$

$$\leq \sum_{x=0}^{n}\binom{n}{x}\sum_{w=0}^{\tilde{x}}e^{-sa_w}q^{-\tilde{x}+w}\exp\left\{a_w\left[\frac{q^2+1}{q^2-1}\frac{n}{2}\left(\log\frac{2\tilde{x}}{\bar{w}}+1\right) + \frac{(q^2+1)^3(q-1)^2}{(q^2-1)^3}\frac{n\pi^2(n-w)}{6(n-1)}\right]\right\}. \tag{E75}$$

We are now in a position to nearly complete the proof. We choose

$$s^* = \frac{1}{2}\frac{q^2+1}{q^2-1}n\log(n) + cn, \tag{E76}$$

where

$$c = \frac{q^2+1}{2(q^2-1)} + \frac{(q^2+1)^3(q-1)^2}{(q^2-1)^3}\frac{\pi^2}{6} + \frac{q^2+1}{(q-1)^2}\left\{\log\left[\frac{320(q-1)(q^n+1)}{9q^n}\right] + \frac{q^2}{q^2-1} + 4\log(q)\right\}. \tag{E77}$$

The $O(n\log(n))$ term in s^* is necessary to cancel the $O(n\log(2\tilde{x}/w))$ term in Eq. (E75). Meanwhile, the first two terms of c are needed to cancel the remaining terms on the right-hand side of Eq. (E75). The next two terms in c are used to cancel

factors on the left-hand side of Eq. (E75) and, finally, the last term is vital for canceling the q^w factor, as follows:

$$
\frac{q^n + 1}{q^n}(q+1)^n(Z - Z_H) \leq \frac{q-1}{q} \sum_{x=0}^{n} \binom{n}{x} \sum_{w=0}^{\tilde{x}} e^{-(s-s^*)a_w} q^{-\tilde{x}+w} \exp\left[-\frac{4w(n-w)}{n-1}\log(q)\right]
$$

$$
\leq \frac{q-1}{q} \sum_{x=0}^{n} \binom{n}{x} \sum_{w=0}^{\tilde{x}} e^{-(s-s^*)a_w} q^{-\tilde{x}-w}
$$

$$
\leq e^{-(s-s^*)a_1} \frac{q-1}{q} \sum_{x=0}^{n} \binom{n}{x} q^{-\tilde{x}} \sum_{w=0}^{\tilde{x}} q^{-w} \tag{E78}
$$

$$
\leq e^{-(s-s^*)a_1} \sum_{x=0}^{n} \binom{n}{x} q^{-\tilde{x}}
$$

$$
\leq e^{-(s-s^*)a_1} \sum_{x=0}^{n} \binom{n}{x} \left(q^{-x} + q^{-n+x}\right)
$$

$$
= 2e^{-(s-s^*)a_1} \left(\frac{q+1}{q}\right)^n.
$$

It is in Eq. (E78), where we use $a_1 \leq a_w$ to pull the exponential out from the sum, that the assumption $s \geq s^*$ is necessary. This is the only place where it is needed. As $Z_H = 2/(q^n + 1)$, we then have

$$
Z \leq Z_H[1 + e^{-(s-s^*)a_1}], \tag{E79}
$$

where $a_1 = (q-1)^2/[n(q^2+1)]$. Defining $a = na_1/2$, this completes the proof of the upper bound. ∎

4. Lower-bound proof

Theorem 12 (Theorem 6 from main text). *For the complete-graph architecture of size s on n qudits with local dimension q, the collision probability satisfies*

$$
Z \geq \frac{Z_H}{2} \exp\left[\frac{\log(q)}{q+1} \exp\left\{\log(n) + s\log\left[1 - \frac{2(q^2-1)}{n(q^2+1)}\right]\right\}\right]. \tag{E80}
$$

Corollary 6. *For the complete-graph architecture, let s_{AC} be the minimum circuit size, as a function of n, such that*

$$
Z \leq 2Z_H. \tag{E81}
$$

Then it must hold that

$$
\left| s_{AC} - \frac{q^2+1}{2(q^2-1)} n\log(n) \right| = O(n). \tag{E82}
$$

Proof. The upper bound on Z in Theorem 11 implies that

$$
s_{AC} \leq \frac{q^2+1}{2(q^2-1)} n\log(n) + O(n). \tag{E83}
$$

Meanwhile, since $s = s_{AC}$ implies $Z \leq 2Z_H$, the bound in Theorem 12 implies that

$$
\exp\left[\frac{\log(q)}{q+1} \exp\left\{\log(n) + s_{AC}\log\left[1 - \frac{2(q^2-1)}{n(q^2+1)}\right]\right\}\right] \leq 4 \tag{E84}
$$

and thus

$$s_{AC} \geq \frac{\log(n) - \log\left[\frac{(q+1)\log(4)}{\log(q)}\right]}{-\log\left[1 - \frac{2(q^2-1)}{n(q^2+1)}\right]}$$

$$\geq \left\{\log(n) - \log\left[\frac{(q+1)\log(4)}{\log(q)}\right]\right\}\left[\frac{n(q^2+1)}{2(q^2-1)} - 1\right]$$

$$= \frac{n(q^2+1)}{2(q^2-1)}\log(n) - O(n), \tag{E85}$$

where we use the general inequality $-1/\log(1-u) \geq 1/u - 1$. ∎

Proof of Theorem 12. The structure of the proof is very similar to Theorem 8 for general architectures. We use the framework of the biased random walk.

Let $x := \gamma^{(t)}$. The transition rule is such that

$$\gamma^{(t+1)} = \begin{cases} x, & \text{with probability } 1 - \dfrac{2x(n-x)}{n(n-1)}, \\ x-1, & \text{with probability } \dfrac{2x(n-x)}{n(n-1)}\dfrac{q^2}{q^2+1}, \\ x+1, & \text{with probability } \dfrac{2x(n-x)}{n(n-1)}\dfrac{1}{q^2+1}, \end{cases} \tag{E86}$$

and so

$$\underset{P_b}{\mathbb{E}}[\gamma^{(t+1)}|\gamma^{(t)} = x] = x - \frac{2x(n-x)}{n(n-1)}\frac{q^2-1}{q^2+1}$$

$$\geq x\left[1 - \frac{2(q^2-1)}{n(q^2+1)}\right]. \tag{E87}$$

As this is true for all x, when we have some probability distribution Λ over values of x, it still holds that

$$\underset{P_b}{\mathbb{E}}[\gamma^{(t+1)}] = \sum_{x=0}^{n}\underset{\Lambda}{\Pr}[\gamma^{(t)} = x]\underset{P_b}{\mathbb{E}}[\gamma^{(t+1)}|\gamma^{(t)} = x]$$

$$\geq \sum_{x=0}^{n}\underset{\Lambda}{\Pr}[\gamma^{(t)} = x]x\left[1 - \frac{2(q^2+1)}{n(q^2-1)}\right]$$

$$= \left[1 - \frac{2(q^2+1)}{n(q^2-1)}\right]\sum_{x=0}^{n}\underset{\Lambda}{\Pr}[\gamma^{(t)} = x]x$$

$$= \left[1 - \frac{2(q^2-1)}{n(q^2+1)}\right]\underset{\Lambda}{\mathbb{E}}[\gamma^{(t)}] \tag{E88}$$

and by applying this equation recursively from the starting distribution Λ_b, we find that

$$\underset{P_b,\Lambda_b}{\mathbb{E}}[\gamma^{(s)}] \geq \left[1 - \frac{2(q^2-1)}{n(q^2+1)}\right]^s\underset{\Lambda_b}{\mathbb{E}}[\gamma^{(0)}]$$

$$= \left[1 - \frac{2(q^2-1)}{n(q^2+1)}\right]^s\frac{n}{q+1}. \tag{E89}$$

By convexity, we have $\mathbb{E}[q^x] \geq q^{\mathbb{E}[x]}$, and hence

$$
\begin{aligned}
Z &= \frac{1}{q^n} \mathop{\mathbb{E}}_{P_b, \Lambda_b} [q^{|\vec{\gamma}^{(s)}|}] \\
&\geq \frac{1}{q^n} \exp \left\{ \log(q) \frac{n}{q+1} \left[1 - \frac{2(q^2-1)}{n(q^2+1)} \right]^s \right\} \\
&\geq \frac{Z_H}{2} \exp \left[\frac{\log(q)}{q+1} \exp \left\{ \log(n) + s \log \left[1 - \frac{2(q^2-1)}{n(q^2+1)} \right] \right\} \right].
\end{aligned}
\tag{E90}
$$

■

5. Delayed proofs of lemmas

Proof of Lemma 6. The first equation follows fairly straightforwardly from Lemma 1. Note that the factor in parentheses on the left-hand side is $\bar{q}_{v,a}/(\bar{q}_{v,a}^2 + 1)$ as defined in Definition 6, as well as the fact that Λ_v contains walks that start at v and end at $v - 1$ or $n - v + 1$. The walks that start at v and end at $v - 1$ are covered by Lemma 1 with $x \to v, y \to v - 1$, and $m \to n - 2v + 2$. The walks that start at v and end at $n - v + 1$ are equivalent to walks starting at $n - v$ and ending at $v - 1$ and are thus covered by Lemma 1 with $x \to n - v, y \to v - 1$, and $m \to n - 2v + 2$. By summing the results from these two substitutions, we obtain the quantity

$$
\frac{1}{1 - \bar{q}_{v,a}^{-2(n-2v+2)}} \left(\bar{q}_{v,a}^{-1} - \bar{q}_{v,a}^{-2n+4v-3} \right) + \frac{1}{1 - \bar{q}_{v,a}^{-2(n-2v+2)}} \left(\bar{q}_{v,a}^{-n+2v-1} - \bar{q}_{v,a}^{-n+v-3} \right) = \bar{q}_{v,a}^{-1} \frac{1 + \bar{q}_{v,a}^{-n+2v}}{1 + \bar{q}_{v,a}^{-n+2v-2}},
\tag{E91}
$$

which proves that the first equation in the lemma is correct.

The second equation is not a direct application of Lemma 1 but it can be shown by a similar method. Define

$$
\Xi_{x,v} = \{ \psi \in \Psi_x : v \leq m(\psi), M(\psi) \leq n - v \},
\tag{E92}
$$

so that $\Xi_v = \Xi_{v,v}$. Moreover, for fixed v, let

$$
I(x) := \sum_{\alpha \in \Xi_{x,v}} \left[\frac{q}{q^2+1} \frac{1}{1 - \lambda_v(1 - e^{-a})} \right]^{L[\alpha]} = \sum_{\alpha \in \Xi_{x,v}} \left(\frac{\bar{q}_{v,a}}{\bar{q}_{v,a}^2 + 1} \right)^{L[\alpha]}.
\tag{E93}
$$

The function $I(x)$ obeys the recursion relation

$$
I(x) = 1 + \frac{\bar{q}_{v,a}}{\bar{q}_{v,a}^2 + 1} [I(x-1) + I(x+1)],
\tag{E94}
$$

since there is one term in the sum corresponding to the length-0 trajectory, which contributes 1, but all other terms appear either in $I(x-1)$ or $I(x+1)$ reduced by factor $\bar{q}_{v,a}/(\bar{q}_{v,a}^2 + 1)$. The general solution to this recursion relation is

$$
I(x) = \frac{\bar{q}_{v,a}^2 + 1}{(\bar{q}_{v,a} - 1)^2} + A\bar{q}_{v,a}^x + B\bar{q}_{v,a}^{-x}
\tag{E95}
$$

for some constants A and B. Here is where we rely on boundary conditions. We must have $I(v-1) = I(n-v+1) = 0$, since these sums do not include any terms. This allows us to solve for A and B and find

$$A = -\frac{\bar{q}_{v,a}^2 + 1}{(\bar{q}_{v,a} - 1)^2} \frac{\bar{q}_{v,a}^{-n+v-1}}{1 + \bar{q}_{v,a}^{-n+2v-2}}, \tag{E96}$$

$$B = -\frac{\bar{q}_{v,a}^2 + 1}{(\bar{q}_{v,a} - 1)^2} \frac{\bar{q}_{v,a}^{v-1}}{1 + \bar{q}_{v,a}^{-n+2v-2}}, \tag{E97}$$

$$I(v) = \frac{\bar{q}_{v,a}^2 + 1}{(\bar{q}_{v,a} - 1)^2} \left(1 - \bar{q}_{v,a}^{-1} \frac{1 + \bar{q}_{v,a}^{-n+2v}}{1 + \bar{q}_{v,a}^{-n+2v-2}} \right). \tag{E98}$$

∎

Proof of Lemma 7. Define $\eta = \lambda_x (1 - e^{-a})$ and let $\zeta = 1 - (1 - \eta)^2$. Thus, $1 - \eta = \sqrt{1 - \zeta}$.

We can write

$$\bar{q}_{x,a}^{-1} = q^{-1} \left(\frac{q^2 + 1}{2} \right) \left(\sqrt{1 - \zeta} \right) \left[1 - \sqrt{1 - \frac{4q^2}{(q^2 + 1)^2 (1 - \zeta)}} \right]$$

$$= q^{-1} \left(\frac{q^2 + 1}{2} \right) \left[\sqrt{1 - \zeta} - \frac{q^2 - 1}{q^2 + 1} \sqrt{1 - \frac{(q^2 + 1)^2}{(q^2 - 1)^2} \zeta} \right]$$

$$\leq q^{-1} \left(\frac{q^2 + 1}{2} \right) \left\{ 1 - \frac{1}{2}\zeta - \frac{q^2 - 1}{q^2 + 1} \left[1 - \frac{(q^2 + 1)^2}{2(q^2 - 1)^2} \zeta - \frac{(q^2 + 1)^4}{2(q^2 - 1)^4} \zeta^2 \right] \right\}$$

$$= q^{-1} \left[1 + \frac{1}{2} \frac{q^2 + 1}{q^2 - 1} \zeta + \frac{(q^2 + 1)^4}{4(q^2 - 1)^3} \zeta^2 \right]$$

$$\leq q^{-1} \exp \left[\frac{1}{2} \frac{q^2 + 1}{q^2 - 1} \zeta + \frac{(q^2 + 1)^4}{4(q^2 - 1)^3} \zeta^2 \right]$$

$$\leq q^{-1} \exp \left[\frac{q^2 + 1}{q^2 - 1} a\lambda_x + \frac{(q^2 + 1)^4}{(q^2 - 1)^3} a\eta\lambda_x \right], \tag{E99}$$

which is equal to the lemma statement, where in the first inequality we utilize $1 - (u/2) - (u^2/2) \leq \sqrt{1 - u} \leq 1 - (u/2)$, in the second inequality we use $1 + u \leq \exp(u)$, and in the third inequality we use $\zeta \leq 2\eta \leq 2a\lambda_x$. The condition on a is necessary to ensure that $\bar{q}_x$ is real. ∎

APPENDIX F: APPROXIMATE 2-DESIGNS AND ANTICONCENTRATION

In this appendix, we clarify the relation between approximate unitary 2-designs and anticoncentration. As we discussed in the text, forming a unitary 2-design is a sufficient condition for anticoncentration.

First, we recall some definitions. The *k*-fold channel of an operator $\mathcal{O}$ with respect to a probability distribution μ on the unitary group $\mathcal{U}(q^n)$ is defined as

$$\Phi_\mu^{(k)}(\mathcal{O}) := \int d\mu(U) \, U^{\otimes k}(\mathcal{O})U^{\dagger \otimes k}. \tag{F1}$$

We denote the channel with respect to the Haar measure on the unitary group as $\Phi_H^{(k)}$. The diamond norm of a quantum channel Φ is defined as $\|\Phi\|_\diamond := \sup_{\psi, D} \|\Phi \otimes \mathcal{I}_D(\psi)\|_1$, where $\mathcal{I}_D$ is the identity channel on a D-dimensional ancilla and ψ is a state on the entire system.

Definition 7 (Approximate designs). *A probability distribution μ on $\mathcal{U}(q^n)$ is an ε-approximate unitary k-design if the k-fold channels obey*

$$\left\| \Phi_\mu^{(k)} - \Phi_H^{(k)} \right\|_\diamond \leq \varepsilon. \tag{F2}$$

For a given k, if $\varepsilon = 0$, we say that the distribution forms an exact k-design.

A weaker notion of approximate design involves the operator norm of the moment operators, sometimes referred to as the TPE condition. The vectorization isomorphism uniquely maps channels to operators, with which we can define the kth moment operator from the k-fold channel for a probability distribution μ on the unitary group $\mathcal{U}(q^n)$ as

$$\widehat{\Phi}_\mu^{(k)} := \text{vec}\left(\Phi_\mu^{(k)}\right) = \int d\mu(U)\, U^{\otimes k} \otimes U^{*\otimes k}. \tag{F3}$$

For convenience, we denote $U^{\otimes k,k} := U^{\otimes k} \otimes U^{*\otimes k}$.

Definition 8 (Weak approximate designs). *A probability distribution μ on $\mathcal{U}(q^n)$ is a weak ε-approximate unitary k-design if the kth moment operators obey*

$$\left\| \widehat{\Phi}_\mu^{(k)} - \widehat{\Phi}_H^{(k)} \right\|_\infty \leq \varepsilon. \tag{F4}$$

The expectation of the collision probability for completely Haar-random unitaries is $Z_H = \mathbb{E}_H[Z] = 2/(q^n + 1) \leq 2/q^n$ and thus anticoncentrates with $\alpha = 1/2$ as defined in Definition 4. But as the collision probability is a second-moment quantity, where $p_U(x)^2 = |\langle x| U |1^n\rangle|^4$, for an exact unitary 2-design μ, we find

$$Z = \mathbb{E}_\mu\left[\sum_x p_U(x)^2 \right] = \mathbb{E}_H\left[\sum_x p_U(x)^2 \right] = \frac{2}{q^n + 1} \tag{F5}$$

and thus also $1/2$-anticoncentrates, where $\mathbb{E}_H[\cdots]$ denotes the expectation with respect to the Haar measure on the unitary group.

Proposition 1. *An ε-approximate 2-design μ with $\varepsilon = 1/q^{2n}$ has a collision probability of $Z = \mathbb{E}_\mu[\sum_x p_U(x)^2] \leq 3/q^n$ and is thus a $1/3$ anticoncentrator. Moreover, the same holds for a weak ε-approximate 2-design (TPE) μ with $\varepsilon = 1/q^{2n}$.*

Proof. For an ε-approximate 2-design in the diamond norm, we find

$$\mathbb{E}_\mu\left[p_U(x)^2\right] = \mathbb{E}_\mu\left[|\langle x|U|1^n\rangle|^4\right] - \mathbb{E}_H\left[|\langle x|U|1^n\rangle|^4\right] + \mathbb{E}_H\left[|\langle x|U|1^n\rangle|^4\right]$$

$$= \text{Tr}\left(|x\rangle\langle x|^{\otimes 2} \left\{ \mathbb{E}_\mu\left[U^{\otimes 2}(|1^n\rangle\langle 1^n|)U^{\dagger\otimes 2} \right] - \mathbb{E}_H\left[U^{\otimes 2}(|1^n\rangle\langle 1^n|)U^{\dagger\otimes 2} \right] \right\} \right) + \frac{2}{q^n(q^n + 1)}$$

$$\leq \left\| |x\rangle\langle x|^{\otimes 2}\left[\Phi_\mu^{(2)}(|1^n\rangle\langle 1^n|) - \Phi_H^{(2)}(|1^n\rangle\langle 1^n|) \right] \right\|_1 + \frac{2}{q^n(q^n + 1)}$$

$$\leq \left\| |x\rangle\langle x|^{\otimes 2} \right\|_\infty \left\| [\Phi_\mu^{(2)} - \Phi_H^{(2)}](|1^n\rangle\langle 1^n|) \right\|_1 + \frac{2}{q^n(q^n + 1)}$$

$$\leq \frac{2}{q^n(q^n + 1)} + \varepsilon, \tag{F6}$$

where we write the difference in terms of the 2-fold channels, in the second to last line use Hölder's inequality, and in the last line use the definition of the diamond norm and the definition of an ε-approximate 2-design.

Given the definition of an approximate design in terms of the diamond norm, we must take the error to be exponentially small. Thus, for an approximate 2-design μ with $\varepsilon = 1/q^{2n}$, the collision probability is $Z \leq 3/q^n$ and thus $1/q^{2n}$-approximate unitary 2-designs in diamond norm anticoncentrate with $\alpha = 1/3$.

For a weak ε-approximate 2-design in the operator norm (TPE), we proceed similarly,

$$\mathbb{E}_\mu\big[p_U(x)^2\big] = \mathbb{E}_\mu\big[|\langle x|U|1^n\rangle|^4\big] - \mathbb{E}_H\big[|\langle x|U|1^n\rangle|^4\big] + \mathbb{E}_H\big[|\langle x|U|1^n\rangle|^4\big]$$

$$= \text{Tr}\Big(|1^n x\rangle\langle 1^n x|^{\otimes 2}\big(\mathbb{E}_\mu[U^{\otimes 2,2}] - \mathbb{E}_H[U^{\otimes 2,2}]\big)\Big) + \frac{2}{q^n(q^n+1)}$$

$$\leq \Big\|\widehat{\Phi}_\mu^{(2)} - \widehat{\Phi}_H^{(2)}\Big\|_\infty + \frac{2}{q^n(q^n+1)}$$

$$\leq \frac{2}{q^n(q^n+1)} + \varepsilon, \tag{F7}$$

where we write the difference in terms of the 2-fold moment operators, in the second to last line use Hölder's inequality, and in the last line use the definition of a weak ε-approximate 2-design. Again, we must take the error to be exponentially small. For $\varepsilon = 1/q^{2n}$, the collision probability is $Z \leq 3/q^n$ and thus $1/q^{2n}$-approximate unitary 2-designs in operator norm anticoncentrate with $\alpha = 1/3$. ∎

As n-qudit RQCs on the 1D architecture are known to form ε-approximate unitary 2-designs in $O(n + \log(1/\varepsilon))$ depth [15,49], anticoncentration for 1D random circuits in linear depth is an immediate corollary. Moreover, an n-independent upper bound on the spectral gap for the 1D architecture [15] implies that they form weak approximate 2-designs in $O(\log(1/\varepsilon))$ depth. By Proposition 1, where we must take $\varepsilon = 1/q^{2n}$, this again requires linear depth for 1D RQCs.

For nonlocal RQCs on the complete-graph architecture, the best known upper bounds on the 2-design circuit size are $O(n^2)$ [13]. However, it has been conjectured that this can be improved to $O(n\log(n))$, in which case anticoncentration and 2-designs could occur at the same depth for the complete-graph circuit architecture.

To argue that anticoncentration must be distinct from the 2-design property, we consider lower bounds on the 2-design depth for RQCs on the 1D architecture. The spectral gap of the second moment of a probability distribution ν on the unitary group is defined as $g(\nu) := \|\widehat{\Phi}_\nu^{(2)} - \widehat{\Phi}_H^{(2)}\|_\infty$. An n-independent bound on the spectral gap for 1D RQCs has been proved in Ref. [15]. This implies that the behavior of the spectral gap for 1D RQCs of depth d must be $g(\nu_{\text{1D RQC}}) = (1 - 1/c)^d$, for some constant $c > 1$. Further recalling that the operator norm can be written as $\|M\|_\infty = \max_y\langle y|M|y\rangle$, this implies that some states requires linear depth in order to become small. Specifically, there is some state $|y\rangle$ on the 4-fold space that requires the 1D circuit depth to be at least $d = \Omega(n)$ in order for the second-moment operator for 1D RQCs $\mathbb{E}_U\big[\langle y|U^{\otimes 2,2}|y\rangle\big]$ to approach the minimal Haar value.

[1] A. Nahum, J. Ruhman, S. Vijay, and J. Haah, Quantum Entanglement Growth under Random Unitary Dynamics, Phys. Rev. X **7**, 031016 (2017).

[2] C. von Keyserlingk, T. Rakovszky, F. Pollmann, and S. Sondhi, Operator Hydrodynamics, OTOCs, and Entanglement Growth in Systems without Conservation Laws, Phys. Rev. X **8**, 021013 (2018).

[3] A. Nahum, S. Vijay, and J. Haah, Operator Spreading in Random Unitary Circuits, Phys. Rev. X **8**, 021014 (2018).

[4] P. Hayden and J. Preskill, Black holes as mirrors: Quantum information in random subsystems, JHEP **09**, 120 (2007).

[5] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, Nat. Phys. **14**, 595 (2018).

[6] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. Brandão, and D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).

[7] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, and D. Fan *et al.*, Strong quantum computational advantage using a superconducting quantum processor, arXiv:2106.14734 [quant-ph].

[8] O. C. O. Dahlsten, R. Oliveira, and M. B. Plenio, The emergence of typical entanglement in two-party random processes, J. Phys. A **40**, 8081 (2007).

[9] R. Oliveira, O. C. O. Dahlsten, and M. B. Plenio, Generic Entanglement Can Be Generated Efficiently, Phys. Rev. Lett. **98**, 130502 (2007).

[10] W. Brown and O. Fawzi, Scrambling speed of random quantum circuits, arXiv:1210.6644 [quant-ph].

[11] W. Brown and O. Fawzi, Decoupling with random quantum circuits, Comm. Math. Phys. **340**, 867 (2015).

[12] W. Brown and O. Fawzi, in *IEEE International Symposium on Information Theory—Proceedings* (2013), p. 346, arXiv:1312.7646 [quant-ph].

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

[14] W. G. Brown and L. Viola, Convergence Rates for Arbitrary Statistical Moments of Random Quantum Circuits, Phys. Rev. Lett. **104**, 250501 (2010).

[15] F. G. S. L. Brandão, A. W. Harrow, and M. Horodecki, Local random quantum circuits are approximate polynomial-designs, Commun. Math. Phys. **346**, 397 (2016).

[16] A. Harrow and S. Mehraban, Approximate unitary *t*-designs by short random quantum circuits using nearest-neighbor and long-range gates, arXiv:1809.06957 [quant-ph].

[17] N. Hunter-Jones, Unitary designs from statistical mechanics in random quantum circuits, arXiv:1905.12053 [quant-ph].

[18] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, On the complexity and verification of quantum random circuit sampling, Nat. Phys. **15**, 159 (2019).

[19] R. Movassagh, Efficient unitary paths and quantum computational supremacy: A proof of average-case hardness of random circuit sampling, arXiv:1810.04681 [quant-ph].

[20] R. Movassagh, Quantum supremacy and random circuits, arXiv:1909.06210 [quant-ph].

[21] A. Bouland, B. Fefferman, Z. Landau, and Y. Liu, Noise and the frontier of quantum supremacy, arXiv:2102.01738 [quant-ph].

[22] Y. Kondo, R. Mori, and R. Movassagh, Improved robustness of quantum supremacy for random circuit sampling, arXiv:2102.01960 [quant-ph].

[23] S. Aaronson and A. Arkhipov, in *Proceedings of the Annual ACM Symposium on Theory of Computing* (2011). arXiv:1011.3245 [quant-ph].

[24] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Average-Case Complexity versus Approximate Simulation of Commuting Quantum Computations, Phys. Rev. Lett. **117**, 080501 (2016).

[25] T. Morimae, Hardness of classically sampling the one-clean-qubit model with constant total variation distance error, Phys. Rev. A **96**, 040302 (2017).

[26] A. Bouland, J. F. Fitzsimons, and D. E. Koh, in *Proceedings of the 33rd Computational Complexity Conference* (2018), p. 21:1, arXiv:1709.01805 [quant-ph].

[27] D. Hangleiter, J. Bermejo-Vega, M. Schwarz, and J. Eisert, Anticoncentration theorems for schemes showing a quantum speedup, Quantum **2**, 65 (2018).

[28] A. M. Dalzell, A. W. Harrow, D. Enshan Koh, and R. L. La Placa, How many qubits are needed for quantum computational supremacy? Quantum **4**, 264 (2020).

[29] T. Morimae and S. Tamaki, Additive-error fine-grained quantum supremacy, Quantum **4**, 329 (Sept., 2020).

[30] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Achieving quantum supremacy with sparse and noisy commuting quantum computations, Quantum **1**, 8 (2017).

[31] X. Gao and L. Duan, Efficient classical simulation of noisy quantum computation, arXiv:1810.03176 [quant-ph].

[32] B. Barak, C.-N. Chou, and X. Gao, Spoofing Linear Cross-Entropy Benchmarking in Shallow Quantum Circuits, arXiv:2005.02421 [quant-ph].

[33] J. Haferkamp, D. Hangleiter, A. Bouland, B. Fefferman, J. Eisert, and J. Bermejo-Vega, Closing gaps of a quantum advantage with short-time Hamiltonian dynamics, arXiv:1908.08069 [quant-ph].

[34] In the Porter-Thomas distribution, the frequency at which $p_U(x) = p$ is proportional to $\exp(-p/q^{-n})$, illustrated roughly in the middle diagram of Fig. 1.

[35] Note that while Eq. (3) is an ingredient in these arguments, it is not alone sufficient to imply hardness of simulation, as the arguments generally rely on additional unproven conjectures.

[36] E. Onorati, O. Buerschaper, M. Kliesch, W. Brown, A. H. Werner, and J. Eisert, Mixing properties of stochastic quantum Hamiltonians, Commun. Math. Phys. **355**, 905 (2017).

[37] H. Gharibyan, M. Hanada, S. H. Shenker, and M. Tezuka, Onset of random matrix behavior in scrambling systems, JHEP **07**, 124 (2018), [Erratum: JHEP 02, 197 (2019)].

[38] N. Hunter-Jones, Operator growth in random quantum circuits with symmetry, arXiv:1812.08219 [quant-ph].

[39] P. Hayden, S. Nezami, X.-L. Qi, N. Thomas, M. Walter, and Z. Yang, Holographic duality from random tensor networks, JHEP **11**, 009 (2016).

[40] T. Zhou and A. Nahum, Emergent statistical mechanics of entanglement in random unitary circuits, Phys. Rev. B **99**, 174205 (2019).

[41] Y. Bao, S. Choi, and E. Altman, Theory of the phase transition in random unitary circuits with measurements, Phys. Rev. B **101**, 104301 (2020).

[42] C. M. Jian, Y. Z. You, R. Vasseur, and A. W. Ludwig, Measurement-induced criticality in random quantum circuits, Phys. Rev. B **101**, 104302 (2020).

[43] J. Napp, R. L. La Placa, A. M. Dalzell, F. G. S. L. Brandao, and A. W. Harrow, Efficient classical simulation of random shallow 2D quantum circuits, arXiv:2001.00021 [quant-ph].

[44] J. Lopez-Piqueres, B. Ware, and R. Vasseur, Mean-field entanglement transitions in random tree tensor networks, Phys. Rev. B **102**, 064202 (2020).

[45] For example, given an *n*-bit efficiently computable Boolean function f, consider the following circuit U. First, perform a layer of Hadamard gates on every qubit; then perform the $2^n \times 2^n$ diagonal unitary operation $\sum_x (-1)^{f(x)} |x\rangle\langle x|$; then perform another layer of Hadamard gates. It is straightforward to show that $p_U(0^n)$ is proportional to $\text{gap}(f)^2$, where $\text{gap}(f) = |\{x : f(x) = 0\}| - |\{x : f(x) = 1\}|$, which is an extremely difficult quantity to compute classically; it is expected that there exist functions f where the best classical algorithm is essentially a brute-force enumeration over all 2^n inputs x.

[46] One can construct contrived architectures that do not quickly anticoncentrate by partitioning the qudits into many subsets and only rarely choosing a gate that couples qudits from different subsets. We define a property that we call *regularly connected* (Definition 5) to rule out this kind of situation. We prove that it implies anticoncentration in $O(n^2)$ gates and conjecture that this can be improved to $O(n \log(n))$.

[47] B. M. Terhal and D. P. DiVincenzo, Adaptive quantum computation, constant depth quantum circuits and Arthur-Merlin games, Quantum Inf. Comput. **4**, 134 (2004).

[48] R. Goodman and N. R. Wallach, *Representations and Invariants of the Classical Groups* (Cambridge University Press, Cambridge, 2000).

[49] F. G. S. L. Brandão and M. Horodecki, Exponential quantum speed-ups are generic, Quantum Info. Comput. **13**, 901 (2013).