

Architectures for Quantum Simulation Showing a Quantum Speedup

Juan Bermejo-Vega,^{1,*} Dominik Hangleiter,¹ Martin Schwarz,¹ Robert Raussendorf,² and Jens Eisert¹

¹*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*

²*University of British Columbia, Department of Physics and Astronomy,
Vancouver, British Columbia V6T 1Z1, Canada*

 (Received 8 August 2017; revised manuscript received 20 December 2017; published 9 April 2018)

One of the main aims in the field of quantum simulation is to achieve a quantum speedup, often referred to as “quantum computational supremacy,” referring to the experimental realization of a quantum device that computationally outperforms classical computers. In this work, we show that one can devise versatile and feasible schemes of two-dimensional, dynamical, quantum simulators showing such a quantum speedup, building on intermediate problems involving nonadaptive, measurement-based, quantum computation. In each of the schemes, an initial product state is prepared, potentially involving an element of randomness as in disordered models, followed by a short-time evolution under a basic translationally invariant Hamiltonian with simple nearest-neighbor interactions and a mere sampling measurement in a fixed basis. The correctness of the final-state preparation in each scheme is fully efficiently certifiable. We discuss experimental necessities and possible physical architectures, inspired by platforms of cold atoms in optical lattices and a number of others, as well as specific assumptions that enter the complexity-theoretic arguments. This work shows that benchmark settings exhibiting a quantum speedup may require little control, in contrast to universal quantum computing. Thus, our proposal puts a convincing experimental demonstration of a quantum speedup within reach in the near term.

DOI: [10.1103/PhysRevX.8.021010](https://doi.org/10.1103/PhysRevX.8.021010)

Subject Areas: Quantum Physics

I. INTRODUCTION

Quantum devices promise to solve computational problems efficiently for which no classical efficient algorithm exists. The anticipated device of a universal quantum computer would solve problems for which no efficient classical algorithm is known, such as integer factorization [1] and simulating many-body Hamiltonian dynamics [2]. However, the experimental realization of such a machine requires fault-tolerant protection of universal dynamics against arbitrary errors [3–5][6,7]. However, because of a prohibitive error-correction overheads, this requires qubit numbers beyond reach in available quantum devices. This does not mean, however, that the demonstration of a computational quantum advantage is unfeasible with current technology.

Indeed, in recent years, it has become a major milestone in quantum information processing to identify and build a simple (perhaps nonuniversal) quantum device that offers a

large (exponential or superpolynomial) computational speedup compared to classical supercomputers, disregarding the concrete practical applicability of the solved problem. The demonstration of such an advantage based on solid complexity-theoretic arguments is often referred to as “quantum computational supremacy” [8]. This important near-term goal still constitutes a significant challenge—as technological advances seem to be required to achieve it—as well as significant efforts in theoretical computer science, physics, and the numerical study of quantum many-body systems: after all, intermediate problems have to be identified with the potential to act as vehicles in the demonstration of a quantum advantage, in the presence of realistic errors.

There is already evidence that existing dynamical quantum simulators [9,10] have the ability to outperform classical supercomputers. Specifically, the experiments of Refs. [11–14] using ultracold atoms strongly suggest such a feature: They probe situations in which, for short times [11] or in one spatial dimension [12,13], the system can be classically simulated in a perfectly efficient fashion using tensor network methods and can even be equipped with rigorous error bounds. However, for long times [11] or in higher spatial dimensions [12,13], such a classical simulation is no longer feasible with state-of-the-art simulation tools. Still, taking the role of devil’s advocate, one may argue that this could be a consequence of a lack of imagination, as there could—in principle—be a simple

*Corresponding author.
jbermejovega@gmail.com

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

classical description capturing the observed phenomena. Hence, a complexity-theoretic demonstration of a quantum advantage of quantum simulators outperforming classical machines is highly desirable [15]. Not all physically meaningful quantum simulations can be underpinned by such an argument, but it goes without saying that the field of quantum simulation would be seriously challenged if such a rigorous demonstration was out of reach.

Several settings for achieving a quantum speedup have been proposed [16–20] based on quantum processes that are classically hard to simulate probabilistically unless the polynomial hierarchy (PH) collapses. These processes remain difficult to simulate up to realistic (additive) errors, assuming further plausible complexity-theoretic conjectures. The proof techniques used build upon earlier proposals giving rise to such a collapse [21,22]. However, at the same time, they still come with substantial experimental challenges.

This work constitutes a significant step towards identifying physically realistic settings that show a quantum speedup by laying out a versatile and feasible family of architectures based on quenched, local, many-body dynamics. We remain close to what one commonly conceives as a dynamical quantum simulator [9–11,23], which is set up to probe exciting physics of interacting quantum systems. Indeed, it is our aim to remain as close as possible to experimentally accessible or at least realistic prescriptions, closely reminiscent of dynamical quantum simulators, while at the same not compromising the rigorous complexity-theoretic argument.

Our specific contributions are as follows. We focus on schemes in which random initial states are prepared on the 2D square lattices of suitable periodicity, followed by unitary constant-time dynamics following a quench to a local nearest-neighbor (NN), translation-invariant (TI) Hamiltonian. These are prescriptions that are close to those that can be routinely implemented with cold atoms in optical lattices [9,24–26]. Since evolution time is short, decoherence will be comparably small. In a last step, all qubits are measured in a fixed identical basis, producing an outcome distribution that is hard to classically sample from within constant ℓ_1 -norm error, requiring no postselection. Technically, our results implement sampling over new families of NNTI 2-local constant-depth [21] IQP circuits [17,22,27] (IQP circuits standing for instantaneous quantum polynomial time circuits, instances of commuting quantum circuits). We build upon and develop a type of setting [28] in which resource states for measurement-based quantum computation are prepared (MBQC) [29] but subsequently nonadaptively measured. We lay out the complexity-theoretic assumptions made, detail how they are analogous to those in Refs. [16–20], and present results on anticoncentration.

By doing so, we arrive at surprisingly flexible and simple NNTI quantum simulation schemes on square lattices, requiring different kinds of translational invariance in the

preparation. Interestingly, and possibly counterintuitively, our schemes share the feature that the final state before the readout step can be efficiently and rigorously certified in its correctness. This is further achieved via simple protocols that involve on-site measurements and a number of samples of the resource state that scales quadratically in the system size. The possibility of certification is unique to our approach. In fact, from the quadratically many samples of the prepared state, one can directly and rigorously infer about the very quantity that is used in the complexity-theoretic argument. We believe that the possibility of such certification is crucial when it comes to unambiguously arguing that a quantum device has the potential to show a true quantum speedup.

Based on our analysis, we predict that short-time certifiable quantum-simulation experiments on as little as 50×50 qubit square lattices should be intractable for state-of-the-art classical computers [18,30]. It is important to stress that this assessment includes the rigorous certification part, and no hidden or unknown costs have to be added to this. Our proposed experiments are particularly suited to qubits arranged in two-dimensional lattices, e.g., cold atoms in optical lattices [9,24–26] with qubits encoded in hyperfine levels of atoms. When assessing feasible quantum devices, it is crucial to emphasize that system sizes of the kind discussed here are not larger but generally smaller than what is feasible in present-day architectures [9,11,13,26].

II. BASIC SETUP OF THE QUANTUM SIMULATION SCHEMES

We present a new family of simple physical architectures that cannot be efficiently simulated by classical computers with strong evidence (cf. Theorem 1 below). All of them share the basic feature that they are based on the constant-time evolution (quench) of a NNTI Hamiltonian on a square lattice. Each architecture involves three steps:

- Q1 Preparations. Arrange $N := \mu mn$ qubits side by side on an n -row, m -column square lattice \mathcal{L} , with vertices V and edges E , initialized on a product state

$$|\psi_\beta\rangle = \bigotimes_{i=1}^N (|0\rangle + e^{i\beta_i}|1\rangle), \quad \beta \in \{0, \theta\}^N, \quad (1)$$

for fixed $\theta \in \{(\pi/4), (\pi/8)\}$, which is chosen uniformly or randomly with probability p_β (e.g., as a ground state of a disordered model). We consider standard square primitive cells. In one scheme, we allow each vertex to be equipped with an additional qubit, named the “dangling bond qubit.” For this scheme, $\mu = 2$; otherwise, $\mu = 1$.

- Q2 Couplings. Let the system evolve for constant time $\tau = 1$ under the effect of a NNTI Ising Hamiltonian

$$H := \sum_{(i,j) \in E} J_{i,j} Z_i Z_j - \sum_{i \in V} h_i Z_i. \quad (2)$$

This amounts to what is usually referred to as a quench. Local fields $\{h_i\}_i$ and couplings $\{J_{i,j}\}_{i,j}$ are set to implement a unitary $U := e^{iH}$, giving rise to a final ensemble $\Psi := \{p_\beta, |\Psi_\beta\rangle\}_\beta$, $|\Psi_\beta\rangle := U|\psi_\beta\rangle$.

- Q3 Measurement. Measure primitive-cell qubits on the X basis and (if present) dangling-bond qubits on the Z basis. Since the latter can be traded for a measurement in the X basis by a uniform basis rotation, one can equally well measure all qubits in the same basis.

As we will discuss later, all individual steps have been realized with present technology. Note also that Q2 amounts to a constant-depth quantum circuit [31].

A. Physical desiderata and concrete schemes

Before we present the concrete schemes, let us lay out some desiderata of “feasibility in a physical setting” as motivated by a cold-atom setup. Our schemes are constructed along the lines of these desiderata. We require the implementation of each step Q1–Q3 to be as simple as possible. For preparations, couplings, and measurements, we desire the periodicity, as measured by the 2D periods (k_x, k_y) in the xy axes of the TI symmetry, to be small. Coupling strengths should further not scale with the system size and lie within feasible regimes. Last, we want the final measurement to be translationally invariant.

We now present three concrete quantum architectures of the form Q1–Q3 that live up to the above desiderata. We label them Q1–Q3 and illustrate them in Fig. 1:

- I A disordered (DO) product state is prepared on a squared lattice, followed by a quench with an Ising Hamiltonian with couplings $J_{i,j} = h_i = \pi/4$ —which implements controlled- Z (CZ) gates on edges—and a final measurement in the X basis.
- II The initial state is TI with period 1 in one lattice direction and uniformly random in the other ($\text{TI}_{(1,\infty)}$); couplings and measurements are picked as in I.
- III Qubits are prepared on a dangling-bond square lattice. The initial state is TI with period 1 in all

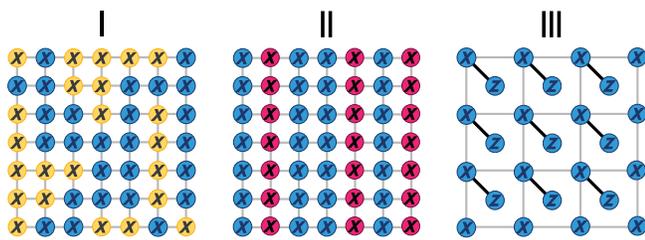


FIG. 1. Architectures I–III. Colors illustrate the rotation angle of the initial state (1): $\beta_i = 0$ (blue), $\beta_i = \pi/4$ (yellow), and $\beta_i = \pi/8$ (crimson). Solid lines between qubits represent Ising-type interactions (2) with coupling constants $J_{i,j} = \pi/4$ (gray) and $J_{i,j} = \pi/8$ (black). X and Z label the basis in which the respective qubits are to be measured.

directions ($\text{TI}_{(1,1)}$). We pick $J_{i,j}, h_i$ as in I and II on bright edges and $J_{i,j} = h_i = \pi/16$ on dark ones—the latter implement controlled- $e^{-i\pi/8Z}$ (CT) gates on dangling bonds. Measurements are in the Z basis for dangling qubits and elsewhere in the X basis.

(See Appendix A for full Hamiltonian specifications.) The resources needed in each architecture are summarized in Table I. It is worth noting that, in all architectures I–III, the state prepared after Q2 is a resource for postselected measurement-based quantum computation postselecting w.r.t. the measurements in Q3 (Sec. VI B) but as such does not amount to universal quantum computation. In fact, architectures I–III require *neither* adaptive measurements (which are key in MBQC [29]) *nor* physical postselection: Our result states that if three plausible complexity-theoretic conjectures hold, a single-shot readout cannot be classically simulated.

III. MAIN RESULT

We now note that, while the above three architectures I–III are physically feasible, the output distributions of measurements cannot be efficiently classically simulated on classical computers, based on plausible assumptions and standard complexity-theoretic arguments.

Theorem 1 (Hardness of classical simulation) If Conjectures 1–3 below are true, then a classical computer cannot sample from the outcome distribution of any architecture I–III up to error $1/22$ in ℓ_1 norm in time $O(\text{poly}(n, m))$.

As in previous works [16,17,20–22,32,33], Theorem 1 relies on plausible complexity-theoretic conjectures. The first, originally adopted in Ref. [21], is a widely believed statement about the structure of an infinite tower of complexity classes known as “the polynomial hierarchy” (PH), the levels of which recursively endow the classes P, NP, and coNP with oracles to previous levels.

Conjecture 1 (Polynomial hierarchy) The polynomial hierarchy is infinite.

The claim generalizes the familiar $P \neq NP$ conjecture in that $P = NP$ would imply a complete collapse of PH to its zeroth level. Furthermore, if two levels $k, k+1$ coincide, then all classes above level k collapse to it. The available evidence for $P \neq NP$ makes Conjecture 1 plausible, for it would be surprising to find a collapse of PH to some level k but not a full one [34] (cf. Ref. [35] for further discussion). Similarly to the Riemann hypothesis in number theory, many theorems in complexity theory have been proven relative to Conjecture 1, probably most notably the Karp-Lipton theorem $NP \not\subseteq P/\text{poly}$ [36].

We highlight that, assuming Conjecture 1 only, a classical computer would still not be able to sample from our experiments either exactly or within any constant relative error (cf. Sec. VI D). However, such a level of accuracy is physically unrealistic because it cannot be

TABLE I. Resource requirements of our hard-to-simulate quantum architectures. Architectures that employ simpler (more ordered) initial states require lattices of higher periodicity and finer controlled rotations. The degree of symmetry of the preparation, evolution, and measurement steps is quantified by the 2D periods indicated by vector subscripts (a, b) . We compare our results to the simplest, previously known, quantum simulation architecture on a planar graph showing a quantum speedup. The underlying complexity-theoretic assumptions needed for these speedups are compared in Sec. III. Above, $X_\theta := e^{-i(\theta/2)Z} X e^{i(\theta/2)Z}$.

Scheme	Geometry	Preparations	Couplings	Measurements
I	Square lattice	DO	$\{\pi/4\}_{(1,1)}$	$\{X\}_{(1,1)}$
II	Square lattice	$\text{TI}_{(1,\infty)}$	$\{\pi/4\}_{(1,1)}$	$\{X\}_{(1,1)}$
III	Dangling-bond square lattice	$\text{TI}_{(1,1)}$	$\{(\pi/4), (\pi/16)\}_{(\sqrt{2},\sqrt{2})}$	$\{X, Z\}_{(\sqrt{2},\sqrt{2})}$
		Previous work		
Reference [19]	Seven-fold brickwork graph	$\text{TI}_{(1,1)}$	$\{\pi/4\}_{(56,2)}$	$\{X, X_{\pm(\pi/4)}, X_{\pm(\pi/8)}\}_{(7,1)}$

achieved by a quantum computer. A goal of this work is to understand how unlikely it is for architectures I–III to be classically intractable under realistic errors.

Our second conjecture, adopted from Ref. [17], is a qubit analog of the “permanent-of-Gaussians” conjecture [16]. It states that partition functions of (unstructured) random Ising models should be equally difficult to approximate in average- and worst-case scenarios. Now, let $(a, b) := (a_1, \dots, a_{N_x}, b_1, \dots, b_{N_z})$ be the outcomes of the X and Z measurements in our architectures, with $b_i = 0$ for I and II. In Appendix B, we show that

$$\text{prob}(a, b|\beta) = |\langle a, b|U|\psi_\beta\rangle|^2 = \frac{|\mathcal{Z}[\pi a, (\pi/4)b + \beta]|^2}{2^{N_x + (N_z/2)}}, \quad (3)$$

where $\mathcal{Z}^{(\alpha, \vartheta)} := \text{tr}(e^{iH^{(\alpha, \vartheta)}})$ is the partition function of a random Ising model on an $n \times m$ square lattice \mathcal{L}_{sq} :

$$H^{(\alpha, \beta)} := \sum_{(i,j) \in E_{\text{sq}}} \frac{\pi}{4} Z_i Z_j - \sum_{i \in V_{\text{sq}}} h_i^{(\alpha, \vartheta)} Z_i, \\ h_i^{(\alpha, \vartheta)} := h_i - \left(\frac{\alpha_i + \vartheta_i}{2} \right), \quad \alpha_i \in \{0, \pi\}, \quad \vartheta_i \in \{0, \theta\}, \quad (4)$$

where $\theta \in \{(\pi/4), (\pi/8)\}$ is chosen as in step Q1 and α (resp. ϑ) is random and DO (resp. either DO or $\text{TI}_{(1,\infty)}$) distributed.

Conjecture 2 [Average-case complexity] For random Ising models as in Eq. (4), approximating $|\mathcal{Z}^{(\alpha, \beta)}|^2$ up to relative error $\frac{1}{4} + o(1)$ for any 0.3 fraction of the field configurations is as hard as in the worst-case scenario.

We complement Conjecture 2 with the following lemma.

Lemma 2 (#P-hardness). Let $H^{(\alpha, \beta)}$ be the Ising model (4) on the $n \times m$ square lattice with either (i) DO-distributed ϑ and $\theta \in \{0, (\pi/4)\}$ or (ii) $\text{TI}_{(1,\infty)}$ -distributed ϑ and $\theta \in \{0, (\pi/8)\}$. Then, for $m \in O(n^2)$, approximating $|\mathcal{Z}^{(\alpha, \beta)}|^2$ with relative error $\frac{1}{4} + o(1)$ is #P-hard.

Thus, accepting Conjecture 2 implies that approximating $|\mathcal{Z}^{(\alpha, \beta)}|^2$ for these models is as hard, on average, as any problem in #P [37]. The proof (Sec. VI C) applies MBQC methods [38] to show that I–III are computationally equivalent to an encoded n -qubit 1D nearest-neighbor circuit comprising random gates of the form

$$\left[\prod_{i=1}^{n-1} CZ_{i,i+1} \right] \left[\prod_{i=1}^n Z_i^{c_i} e^{-i\theta d_i Z_i} H_i \right], \quad c_i, d_i \in \{0, 1\}, \quad (5)$$

where c_i (resp. d_i) is DO (resp. DO-or- $\text{TI}_{(1,\infty)}$) distributed and H is the Hadamard gate. Postselecting such circuits, we can implement two known universal schemes of quantum computation [39,40]. We then exploit the fact that universal quantum-circuit amplitudes are #P-hard to approximate. As a remark, we discuss that the bound $m \in O(n^2)$ in Lemma 2 might not be optimal. In fact, we believe the result should still hold for $m \in O(n)$ (possibly for a different constant error) based on two pieces of evidence.

- (i) On the one hand, our anticoncentration numerics (Appendix C) indicate that $O(n)$ -depth universal random circuits of gates of form (5)—whose output probabilities are in on-to-one correspondence via Eq. (6) with the instances of $|\mathcal{Z}^{(\alpha, \beta)}|^2$ —are Porter-Thomas distributed: The latter is a signature of quantum chaos and of our quantum circuits being approximately Haar random [18,41–45]. Hence, our numerics suggest that our $n \times n$ -qubit lattices efficiently encode chaotic, approximately Haar-random, n -qubit unitaries.
- (ii) On the other hand, we analytically show that #P-hardness arises for $m \in O(n)$ and slightly different choices of input states (resp. dangling bonds) in architectures I and II (resp. III) (cf. Appendix D).

Last, we claim that random circuits of gates of the form (5) anticoncentrate.

Conjecture 3 (Anticoncentration). Let \mathcal{C} be an n -qubit $O(n)$ -depth random circuit of gates of the form (5); then,

$$\text{prob}_x \left(|\langle x|C|0 \rangle|^{\otimes n} \geq \frac{1}{2^n} \right) \geq \frac{1}{e} \quad (6)$$

for a uniformly random choice of $x = (x_1, \dots, x_n)$.

In Sec. VID, we show that Eq. (6) is a sufficient condition for the output distribution of architectures I–III to display anticoncentration. Analogous numerically supported conjectures have been made in Refs. [16,18,46]. Here, we ran exact simulations of random circuits with up to 20 logical qubits to test Conjecture 3 (Appendix C); we observed, first, that the anticoncentration ratio of Eq. (6) quickly converges to $1/e$ with the system size and, second, that measurement outcomes are Porter-Thomas [41] (i.e., exponentially) distributed, which is a signature of chaotic Haar-random unitary processes [18,42–45].

Previously, Refs. [17,20,47] argued that anticoncentration of measurement outcomes on an $n \times n$ lattice should require $\Omega(n)$ physical depth on 2D NN layouts in order not to induce a violation of the counting exponential time hypothesis [48,49]. This contrasts with the constant-depth nature of our proposal. To clarify this discrepancy, we note that our numerical evidence for anticoncentration is for logical n -qubit 1D circuits (5) of depth $O(n)$ (Fig. 6), which implies anticoncentration of the corresponding constant-depth evolution on a lattice of size $n \times O(n)$. Because this encoding introduces a linear overhead factor $O(n)$, there is no contradiction with Refs. [20,47]. More critically, the observed signatures of anticoncentration rule out a potential efficient classical simulation of our schemes via sparse-sampling methods [20,50].

We end this section with a remark: Closest to our work is the approach of Ref. [19], which is also a NNTI nonadaptive MBQC proposal, albeit with larger resource requirements (see Table I for a comparison).

What is more, it requires a stronger hardness assumption with regards to the required level of approximation in that it introduces a variation of Conjecture 2 with a less-natural inverse-exponential additive error (cf. Appendix F). Furthermore, Refs. [18,20,46] gave non-TI schemes based on time-dependent NN random circuits acting on square lattices: The latter approaches require less qubits but also circuits of polynomial depth. In our approach and in that of Ref. [19], circuit depth is traded with ancillas and kept constant, and efficient certification protocols also exist and can be used to determine if the experiment has actually worked, as discussed below.

IV. EFFICIENT CERTIFICATION OF FINAL RESOURCE STATES

It is key to all schemes proposed that the correctness of the final resource-state preparation in the quantum simulation can be efficiently and rigorously certified. Since the prepared state is the ground state of a gapped and frustration-free parent Hamiltonian $H_{\text{parent}} = \sum_i h_i$, Ref. [51] gives a scheme—involving local measurements

only—that certifies the closeness of the prepared state ρ to the anticipated state $|\Psi_\beta\rangle\langle\Psi_\beta|$ immediately before measurement in terms of an upper bound on the trace distance $\| |\Psi_\beta\rangle\langle\Psi_\beta| - \rho \|_1$ [51] (see also Ref. [52]). This directly yields an upper bound on the ℓ_1 -norm distance between the respective measurement outcome distributions.

The key idea of the protocol of Ref. [51] is to estimate the energy $E_\rho = \text{tr}[\rho H_{\text{parent}}]$. This yields a fidelity witness $F(\rho, |\Psi_\beta\rangle\langle\Psi_\beta|) \geq 1 - E_\rho/\Delta$, where Δ is the gap of H_{parent} . This can be done, for example, by measuring the local Hamiltonian terms h_i , which are five- (six-) body observables for architectures I and II (III). Reference [51] showed that this approach requires $O(N^2 \log N)$ samples of the state preparation ρ to estimate a single term $\langle h_i \rangle$ with polynomial accuracy. Hence, the full certification protocol requires $O(N^3 \log(N))$ independent preparations of ρ and five- (six-) body measurements. Though this scaling is efficient, its supra-cubic time scaling and the complexity of the local measurements could render it impractical for near-term experiments with thousands of atoms.

We now introduce three optimizations to the protocol of Ref. [51], analyzed in Appendix E (Lemmas 8–12). First, for any nondegenerate gapped local Hamiltonian with known ground-state energy $E_0 \in O(1)$ and known gap $\Delta \in \Omega(1)$ (but not necessarily frustration free), we show that the sample complexity of the protocol can be reduced to $O(N^2)$ by exploiting parallel measurement sequences of commuting Hamiltonian terms. Second, we show how to implement this more resource-economical protocol, using only on-site measurements, by expanding the latter terms in a local product basis. Third, we introduce a few setting-dependent optimizations for the architectures I–III (Lemma 10), tailored to their underlying square lattice geometry, the explicit tensor product structure of the parent Hamiltonian of their premeasurement states (Appendix E 1), and their translation-invariant symmetry. We emphasize that, as in the sampling measurement step Q3 of our architectures, the optimized certification protocol relies on on-site measurements only.

In the three aforementioned cases (Appendix E), the certification measurement pattern inherits the initial symmetry of the preparation step (Table I), i.e., DO for architecture I, $\text{TI}_{(1,\infty)}$ for architecture II, and $\text{TI}_{(\sqrt{2},\sqrt{2})}$ for architecture III. For architectures I and II, our setting resembles a certification protocol for preparing a family of hypergraph states given in Ref. [53], though states and measurements therein are asymmetric.

The above certification measurement is a slightly more difficult prescription than the experiments as such, yet it is as simple as one could hope since we need to measure additional albeit single-qubit bases. However, it is key to see that the correctness of the final-state preparation of an experiment can be certified even in the absence of a known classical algorithm for sampling its output distribution. This is also in contrast to other similar schemes, where no

efficient rigorous scheme for certification of the final state before measurement is known [16–18,20,46].

Note that the certification protocol of Ref. [51] is stated in terms of noise-free measurements. Nevertheless, for certain noise models, it can be shown that rigorous certification is still possible [19]. Moreover, it is not an unreasonable assumption that local measurements can be benchmarked to a very high precision. Most importantly, the protocol readily accepts imperfect measurements: The imperfect measurements can be seen as perfect measurements, preceded by a quantum channel reflecting the noise. This quantum-channel noise can equally well be seen as acting on the quantum state, reducing the trace-norm closeness to the anticipated state. Hence, certifying ϵ closeness of the state preparation under the assumption of ideal measurements is equivalent to certifying $\epsilon' < \epsilon$ closeness of the state preparation using measurements preceded by a noise channel with \diamond norm at most $\epsilon - \epsilon'$. Setting up a detection scheme living up to the required error bounds (that scales inversely with the system size for each on-site measurement) is demanding but not unrealistic.

Last, we highlight that the certifiability of the final-state preparation is rooted in the fact that the states prepared are ground states of gapped local Hamiltonian models. At the same time, they are injective projected entangled pair states (PEPS) of constant bond dimension [54,55]. The protocols discussed here can hence be seen as PEPS sampling protocols that generate samples from local measurements on PEPS.

Certification protocol.—Let us now outline the precise certification protocol (analyzed in Appendix E) including the required quantum measurements and the postprocessing of the measurement outcomes. We do so in three steps: First, we find the parent Hamiltonians for state preparations $|\Psi_\beta\rangle$ in architectures I–III. Second, we show how on-site measurements are sufficient to obtain a rigorous certificate. Finally, we comment on the reduced sampling complexity $O(N^2)$ of the protocol. We refer the reader to Lemmas 8–12, Appendix E, for proofs of the results described below.

Observing that the resource states $|\Psi_\beta\rangle$ are stabilizer states, all we need to do is find the appropriate stabilizers. The sum of the stabilizers is then a parent Hamiltonian of $|\Psi_\beta\rangle$. For architectures I and II, this yields (cf. Appendix E 1)

$$H_{\text{I,II}} = -\sum_{i \in V} \left(X_{\beta_i, i} \prod_{j: (i,j) \in E} Z_j \right), \quad (7)$$

where $X_{\beta_i, i} = e^{-i(\beta_i/2)Z} X_i e^{i(\beta_i/2)Z}$ is a rotated Pauli- X operator acting on site i and β_i is distributed as described in Q1. Hence, the Hamiltonian consists of N terms that are 5-local except at the boundary of the lattice, where their locality is reduced to 4- or 3-local. In the specific case of architecture

III, a dangling-bond qubit is attached to each qubit via a CT interaction. This yields a two-body term that replaces the $X_{\beta_i, i}$ term in Eq. (7) (see Appendix E 1).

The stabilizers $h_i = X_{\beta_i, i} \prod_{j: (i,j) \in E} Z_j$ in architectures I and II can be measured using on-site measurements in a demolition fashion, by first measuring their tensor components and then multiplying their outcomes using classical postprocessing. (This procedure is reminiscent of the syndrome measurement of subsystem codes [56], and its correctness can easily be seen by decomposing each stabilizer into an eigenbasis. These act on distinct sites and can therefore be measured simultaneously. For the specific case of architecture III, the local Hamiltonian terms have on-site Z factors as well as a two-body component $CTXCT^\dagger$. However, the purpose of measuring these stabilizers in the certification protocol of Ref. [51] (see Appendix E) is to estimate the average energy of the relevant parent Hamiltonian. To this end, we can, without loss of generality (w.l.o.g.), expand $CTXCT^\dagger$ as a sum of product operators [Eq. (E3)], measure the on-site factors appearing in this sum, and infer the target expected value using efficient classical postprocessing. Hence, for our three architectures, on-site measurements suffice.

Finally, the sampling complexity can be reduced from supra-cubic to quadratic $O(N^2)$ by simultaneously measuring commuting stabilizers (directly or by reducing them to on-site measurements as outlined above) on the same state preparation, following a specific pattern (Fig. 2). Precisely, we can define a lattice 2-coloring $V = V_{\text{odd}} \cup V_{\text{even}}$ and simultaneously measure Z on all sites $i \in V_{\text{odd}}$ and $X_{\beta_j, j}$ on every site $j \in V_{\text{even}}$ (or vice versa). Since our Hamiltonian is commuting, each measurement round allows us to sample from the output distribution of $H_{\text{odd}} := \sum_{i \in V_{\text{odd}}} h_i$ ($H_{\text{even}} := \sum_{i \in V_{\text{even}}} h_i$), as shown in Fig. 2. Hence, roughly $N/2$ terms of the form h_i can now be measured in

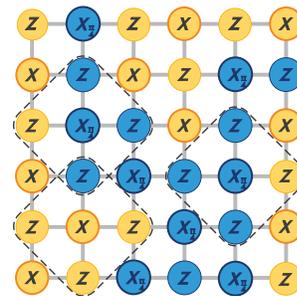


FIG. 2. Certification protocol. We illustrate how our scheme works for architecture I. Qubits with thick (thin) borders denote odd (even) sites in V_{odd} (V_{even}). The figure illustrates a pattern of on-site measurements for one execution of the certification protocol that measures the energy of $H_{\text{odd}} = \sum_{i \in V_{\text{odd}}} h_i$ for the configuration of initial states in Fig. 1. On-site measurements are of type Z , X , and $X_{\pi/4}$. Three Hamiltonian terms whose joint measurement can be simulated from these single-qubit measurements are singled out by the depicted dashed diamonds.

parallel. A simple application of Hoeffding’s bound then shows that we can estimate the expected energy of H_{odd} (H_{even}) using $O(N^2)$ samples. Our proof concludes by noting that $\text{tr}(H\rho) = \text{tr}(H_{\text{odd}}\rho) + \text{tr}(H_{\text{even}}\rho)$.

V. CONCEIVABLE PHYSICAL ARCHITECTURES

We now discuss how the above assumptions are plausible in several physical architectures close to what is available with present technology. What we are considering are large-scale quantum lattice architectures on square lattices \mathcal{L} with a quantum degree of freedom per lattice site. On the level of physical implementation, the most advanced family of such architectures and the most plausible for the anticipated system sizes is that provided by cold atoms in optical lattices [9]. In an optical lattice architecture, internal degrees of freedom are available with hyperfine levels. Also, the encoding in spatial degrees of freedom within double wells is, in principle, conceivable. Large-scale, translationally invariant, controlled- Z interactions—precisely of the type required for the preparation of cluster and graph states [38,57]—are feasible via controlled collisions [25,58]. Actually, the discussion of controlled collisions [58]—and hence the theoretical underpinning of such quenched dynamics—triggered work on cluster states and measurement-based quantum computing and predates this development. Other interactions to nearest neighbors are also conceivable. Interactions such as spin-changing collisions for ^{87}Rb atoms have been experimentally observed [59]. Controlled- T gates require a more sophisticated interaction Hamiltonian. The dangling bonds seem realizable, making use of optical superlattices [11,60]. Single sites—specifically of the sampling type considered here—can be addressed in optical lattice architectures via several methods. In general, quantum-gas microscopes allow for single-site resolved imaging [61,62], even though the type of single-site addressing required here remains a significant challenge. Optical superlattices [11,60] allow us to address entire rows of sites in the same fashion. Sequences of rows where every site is either empty or contains a fixed particle number (following patterns as in Fig. 1 II) can already be routinely prepared [63]. Disordered initial states can also be prepared [12,14,64].

Other architectures are quite conceivable as well. This includes, in particular, large arrays of semiconductor quantum dots allowing for single-site addressing—a setting that has already been employed to simulate the Mott-Hubbard model in the atomic limit [65]—or polaritons or exciton-polariton systems in arrays of microcavities [66]. In this type of architecture, addressing entire rows is also particularly feasible. Superconducting architectures also promise to allow for large-scale array structures of the type anticipated here [67–69]. Trapped ions can also serve as feasible architectures [70]. None of the physical architectures realizes all elements required to the necessary

precision, but at the same time, the prescriptions presented here are comparably close to what can be done.

VI. PROOF OF HARDNESS RESULT

In this section, we prove Lemma 2 and Theorem 1, and develop the main techniques of the paper. The section is organized as follows:

- (i) In Sec. VI A, we use MBQC techniques to develop mappings that allow us to recast architectures I–III as (computationally equivalent) MBQCs on 2D cluster states (as introduced in Ref. [29]).
- (ii) In Sec. VI B, we show that enhancing architectures I–III with the ability (or an oracle) to postselect the outcomes of random variables makes them as powerful as a postselected universal quantum computer (as defined in Ref. [72]).
- (iii) In Sec. VI C, we prove Lemma 2 using earlier findings and a new parallelization technique to implement the 2-local “dense” IQP circuits of Ref. [17] in linear depth on a 1D nearest architecture.
- (iv) In Sec. VI D, we give the proof of our main result, Theorem 1. The proof makes use of Lemma 2 and Stockmeyer’s Theorem [73]. The latter is applied in an analogous way as in the boson-sampling and IQP-circuit settings [16,17] to show that if an efficient classical algorithm can approximately sample from the output distribution of architectures I–III, then an FBPP^{NP} algorithm can approximate a large fraction of the amplitudes in Eq. (3), if the latter are also sufficiently anticoncentrated. By Conjectures 2 and 3, the latter algorithm can solve any problem in P^{#P}, which contains PH via Toda’s theorem [74]. This implies a collapse of the polynomial hierarchy to its third level.

A. Mapping architectures I–III to cluster state MBQCs

We show that any architecture I–III can be mapped via a bijection to a computationally equivalent sequence of X - Y -plane single-qubit measurements on the 2D cluster state [29]. Below, $T := \text{diag}(1, e^{i\pi/4})$ and $\sqrt{T} := \text{diag}(1, e^{i\pi/8})$. First, note that (via teleportation) the effect of measuring a dangling qubit (if present) is equivalent to generating a uniformly random classical bit $b \in \{0, 1\}$ and, subsequently, implementing the gate T^b onto its neighbor; we can thus replace all dangling-bond qubits by introducing a uniformly random measurement of X or $X_{-(\pi/4)} = T^\dagger X T \propto X - Y$ on every primitive qubit. Furthermore, we can rewrite the input $|\psi_\beta\rangle$ in Q1 as

$$|\psi_\beta\rangle = \bigotimes_{i=1}^N \sqrt{T}^{kb_i} |+\rangle^{\otimes N}, \quad (8)$$

where $|+\rangle \propto |0\rangle + |1\rangle$, $k \in \{1, 2\}$ and $b = (b_1, \dots, b_N)$ is a random bit string defined via $b_i := \beta_i/\theta$, with β, θ as in Q1. Since \sqrt{T} gates in Eq. (8) commute with the Hamiltonian (2) and their effect is unobserved by Z measurements, they can be propagated out of the experiment by measuring $X_{-(\pi/8)} = \sqrt{T}^{b_i^\dagger} X_i \sqrt{T}^{b_i}$ instead of X on every primitive qubit $i \in V$. Combining these facts, we obtain the following mappings:

- (C1) Architectures I and III are computationally equivalent to a quantum circuit that prepares a 2D cluster state on their underlying primitive square lattice and measures $\{X, X_{-(\pi/4)}\}$ randomly on each vertex.
- (C2) Architecture II is computationally equivalent to an analogous circuit of random $\{X, X_{-\pi/8}\}$ single-qubit measurements, which chooses the latter measurements to be identical along the columns of the 2D cluster state.

B. Universality of architectures I–III for postselected measurement-based quantum computation

For each of our architectures I–III, we prove that the ensemble $\{p_\beta, |\Psi_\beta\rangle\}_\beta$ is a universal resource for postselected MBQC w.r.t. the measurements in step Q3. Precisely, this means that if the ability to postselect the outcomes of the experiment’s random variables (the qubit outcomes in step Q3 and the random vector β) is provided as an oracle [21,72], then it is possible to implement any poly-size quantum circuit [75] with arbitrarily high fidelity in a subregion of the lattice using (at most) polynomially many qubits. Our proof is constructive and shows how to simulate universal circuits of Clifford + T gates [76] via postselection.

Below, we call a quantum circuit “1D homogeneous” if it consists of 1D nearest-neighbor gates and all nontrivial concurrent operations in a single time step are identical modulo a translation. The latter do not need to be translation invariant, e.g., an arbitrary S -size 1D nearest-neighbor circuit can be serialized to be 1D homogeneous in depth $O(S)$. In Fig. 5, we give an example of an S -size IQP circuit that can be implemented in depth $O(\sqrt{S})$ (by bringing single-qubit gates to the end). The 1D homogeneous circuits, as defined here, can be regarded as examples of quantum cellular automata [77].

Lemma 3 (Postselected universality) Let V be an n -qubit D -depth 1D homogeneous circuit of Clifford + T gates. Then, for any architecture I–III, it is possible to prepare the rightmost primitive qubits of an $O(n) \times O(Dn)$ -qubit lattice on a state $|\psi\rangle := (V|0\rangle^{\otimes n})|0\rangle^{\otimes r}$, $r \in O(n)$, using postselection.

We highlight that the complexity of the simulation in Lemma 3 scales with the depth of the input circuit (not the size), allowing us to parallelize concurrent nearest-neighbor gates. To prove this result, we assume basic knowledge of MBQC on cluster states [29,38]. Additionally, we make use of two technical lemmas.

Lemma 4 (Efficient preparation via MBQC) Let V be an n -qubit D -depth 1D homogeneous Clifford+ T circuit. Then, the state vector $|\psi\rangle := (V|0\rangle^{\otimes n})|0\rangle^{\otimes 3n-2}$ can be efficiently prepared exactly via a MBQC of single-qubit $\{X, X_{\pm\pi/8}\}$ measurements on a $(4n-2) \times O(Dn)$ -qubit 2D cluster state, and even if measurements are constrained to act “quasiperiodically” as follows: For every column, each of its qubits is measured in either the X basis or in one of the $X_{\pm\pi/8}$ bases (where the sign can be picked freely on distinct sites).

Lemma 5 (On-site efficient preparation via MBQC) Let V be an n -qubit D -depth 1D homogeneous Clifford + T circuit. Then, the state vector $|\psi\rangle = V|0\rangle^{\otimes n}$ can be efficiently prepared exactly via a MBQC of single-qubit $\{X, X_{\pm\pi/4}\}$ measurements on an $n \times O(Dn)$ -qubit 2D cluster state.

Lemma 4 is a MBQC implementation of a 1D quantum-computation scheme given in Ref. [39]. Lemma 3 follows from Lemma 3 in Ref. [40] by using the fact that commuting-gate measurement patterns can be applied simultaneously in MBQC.

Proof of Lemma 4.—We first show how to implement a universal set of gates that can be converted to the Clifford + T gate set. We begin by picking a translationally invariant gate set with the desired property [39]

$$\left\{ E := \left(\prod_{i=1}^{M-1} CZ_{i,i+1} \right) \left(\prod_{j=1}^M H_j \right), \quad Y_{\text{all}} := \prod_{j=1}^M Y_j, \right. \\ \left. U_A(\alpha) := \prod_{j=1}^N e^{-i(\alpha/2)A_j}, \quad \text{where } A \in \{X, Z\} \right\}. \quad (9)$$

Above, gates act on a one-dimensional chain of $M := 4n-2$; H is the Hadamard gate; $CZ_{i,i+1}$ is the CZ gate on qubits $i, i+1$; E is a global entangling gate; and E^{M+1} implements a “mirror” permutation $i \rightarrow \bar{i} := M+1-i$ of the qubits. The computation is encoded on n logical qubits with physical positions $[i] := 2i-1, 1 \leq i \leq n$. The remaining qubits are kept in the state $|0\rangle$. Reference [39] shows how to implement generators for the Clifford + T gate set using the following sequences of (9) operations:

- (S1) $\{E^{M+1-i} Y_{\text{all}} E Y_{\text{all}} E^{i-1} U_Z[\mp(\pi/8)]\}$
 $\{E^{M+1-i} Y_{\text{all}} E Y_{\text{all}} E^{i-1} U_Z[\pm(\pi/8)]\}$,
- (S2) $\{E^{M-i} Y_{\text{all}} E Y_{\text{all}} E^i U_X[\mp(\pi/8)]\}$
 $\{E^{M-i} Y_{\text{all}} E Y_{\text{all}} E^i U_X[\pm(\pi/8)]\}$,
- (S3) $\{E^{M-2-[i]} Y_{\text{all}} E Y_{\text{all}} E^{[i]+1} U_X[\mp(\pi/8)] E\}$
 $\{E^{M-2-[i]} Y_{\text{all}} E Y_{\text{all}} E^{[i]+1} U_X[\pm(\pi/8)] E\}$.

Sequence (S1) implements an $e^{\mp i(\pi/8)Z_i}$ gate; (S2), a $e^{\mp i(\pi/8)X_i}$ gate; and (S3), a logical $e^{\mp i(\pi/8)X_{[i]}X_{[i+1]}}$ gate [39]. We now show that any of the above gate sequences can be implemented directly on a MBQC on a $(4n-2) \times O(n)$ -qubit 2D cluster state with quasiperiodic $\{X, X_{\pm\pi/8}\}$ measurements. We make E (resp. nonentangling unitaries), w.l.o.g., act on even steps (resp. odd ones) by introducing identity gates when necessary. We now reorder operations in the creation and measurement of the

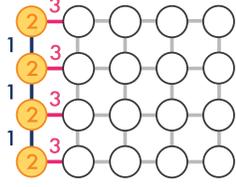


FIG. 3. Mapping the TI quantum computation scheme of Ref. [39] to a cluster-state MBQC.

cluster state as indicated in Fig. 3: Therein, balls denote qubits prepared in $|+\rangle$, steps Fig. 3 (1) and Fig. 3 (2) implement CZ gates for the preparation of the cluster state, and Fig. 3 (3) implements a round of measurements.

In MBQC, Y_{all} gates can be treated as by-product Pauli operators and do not need to be enacted [38]. Furthermore, performing a periodic measurement of X in Fig. 3 (3) implements the E gate in (9). In turn, a quasiperiodic $X_{\pm\pi/8}$ measurement, where observables' signs are chosen adaptively to counteract random by-product operators, can be used to implement E followed by a $U_Z(\pm\pi/8)$ gate. Similarly, we can implement $EU_X(-\pi/8)$ by delaying the measurement to the next step and propagating a $U_Z(-\pi/8)$ backwards: This works because U_Z and U_X never occur in subsequent odd steps in sequences (S1)-(S2)-(S3). We thus have a MBQC simulation of the translation-invariant computation in Ref. [39].

Last, we show that an n -qubit D -depth homogeneous circuit of $e^{\mp i\pi Z_{[i]}/8}$, $e^{\mp i\pi X_{[i]}/8}$, $e^{\mp i\pi X_{[i]}X_{[i+1]}/8}$ gates can be implemented on a $(4n-2) \times O(Dn)$ -qubit cluster-state MBQC using the above protocol. Here, we suggest that measurement patterns on disjoint regions of a cluster-state MBQC can be simultaneously applied for commuting logical gates (hence, also concurrent ones): The latter fact is easily verified in the MBQC's logical-circuit picture [38,78]. \square

Proof of Lemma 5.—Lemma 3 in Ref. [40] shows that performing an $X_{\pm\pi/4}$ measurement on a boundary qubit of an $n \times (n+2)$ one can selectively implement any logical gate of the form $e^{\mp i\pi Z_i/8}$, $e^{\mp i\pi X_i/8}$, $e^{\mp i\pi Z_j X_{j+1}/8}$, $e^{\mp i\pi X_j Z_{j+1}/8}$, $1 \leq i \leq n$, $1 \leq j \leq n-1$ on an n -qubit 1D chain. As in the proof of Lemma 5, measurement patterns associated with commuting gates can be implemented simultaneously. The result of Ref. [40] thus yields an exact $n \times (n+2)$ cluster-state MBQC implementation of any circuit of the form

$$\prod_{i=1}^n e^{-i[(\pi b_i)/8]X_i} C \prod_{i=1}^n e^{-i[(\pi a_i)/8]Z_i}, \quad b_i, a_i \in 0, 1$$

for any n -qubit 1D commuting circuit C of $e^{-i(\pi/8)Z_j X_{j+1}}$, $e^{-i(\pi/8)X_j Z_{j+1}}$ gates. The proof follows the one of Lemma 4. \square

We now proof the main claim of this section.

Proof of Lemma 3.—Recall that our architectures can be recast as a nonadaptive cluster-state MBQC via mappings

(C1) and (C2). Hence, it suffices to show how to prepare $|\psi\rangle = (V|0\rangle^{\otimes n})|0\rangle^{\otimes r}$ exactly for some $r \in O(n)$ using two kinds of operations:

(D1) Postselected random $\{X, X_{-\pi/4}\}$ measurements on an $(n+r) \times O(Dn)$ -qubit cluster state.

(D2) Postselected $\{X, X_{-\pi/8}\}$ measurements, chosen identically (at random) in the vertical (horizontal) direction of an $(n+r) \times O(Dn)$ -qubit cluster state.

Statement (D1) [resp. (D2)] covers the case for I and III (resp. architecture II). To prove (D1) and (D2), we show that if a MBQC scheme on a cluster state is universal w.r.t. a family of X - Y -plane measurements $\{X_{\theta_i}\}_i$, $X_{\theta_i} = e^{-i(\theta_i/2)Z} X e^{i(\theta_i/2)Z}$, then the reduced negative-angle subfamily $\{X_{-\theta_i}\}_i$ is universal for post-MBQC; in combination with Lemmas 4 and 5, the claims follow. Recall that any nonfinal measurement in cluster-state MBQC [29,38] produces a uniformly random outcome $s \in \{0, 1\}$ [cf. Sec. VI C, Eq. (12) for an explicit formula], whose effect in the logical circuit is to introduce a random by-product Pauli operator X^s on its associated qubit line. If not accounted for (e.g., by adapting the measurement basis) and an X_{θ} is subsequently performed, the latter effectively implements an $X_{(-1)^s\theta}$ measurement: This can be seen by propagating X^s forward in the circuit using conjugation relationships, and it is illustrated in Fig. 4. \square

C. #P-hardness of approximating output probabilities (proof of Lemma 2)

In this section, we prove Lemma 2. Our proof below shows that the ability to approximate the given Ising partition function can be used to approximate the output probabilities of the “dense” 2-local long-range IQP circuits of Ref. [17]. The proof further exploits a new technique (Lemma 6) to implement $O(n^2)$ -size long-range IQP circuits in $O(n)$ depth in a 1D nearest-neighbor architecture, which is asymptotically optimal. We regard Lemma 6 of independent interest since the latter dense IQP circuits were argued in Ref. [17] to exhibit a quantum speedup, but, to the best of our knowledge, linear-depth 1D implementations were not previously known. On the other hand, recently, it has been shown that a “sparse” subfamily of the



FIG. 4. (i) Measurement of X_{θ_i} , X_{θ_j} on an edge of a 1D cluster state: The outcomes $s_i, s_j \in \{0, 1\}$ are uniformly random. (ii) The associated logical circuit: The by-product operator X^{s_i} can be propagated forward in the circuit by substituting θ_j with $(-1)^{s_i}\theta_j$. The argument extends to the full cluster by induction, choosing the first qubit to be measured in the X basis (this fixes the input of the logical circuit and does not change its universality properties). The 2D cluster-state case is analogous [29,38].

latter IQP circuits can be implemented in depth $O(n \log n)$ in a 2D nearest-neighbor architecture [20].

1. A 1D linear-depth implementation of dense IQP circuits

We first derive our intermediate result for IQP circuits. For any positive n , we let \mathcal{C} be any “dense” random n -qubit IQP circuit whose gates are uniformly chosen from the set

$$\{e^{i\theta_{i,j}X_iX_j}, e^{i\theta_iX_i} : i, j \in \{1, \dots, n\}, \quad (10)$$

$$\theta_i, \theta_{i,j} \in \left\{ \frac{\pi k}{8}, k = 0, \dots, 7 \right\}, \quad (11)$$

which contains arbitrary long-range interactions in a fully connected architecture.

Lemma 6 (Dense IQP circuits) Dense n -qubit IQP circuits of Eq. (10) gates can be implemented in $\Theta(n)$ depth in a 1D nearest-neighbor architecture.

Proof.—It is easy to see that n^2 -size 2-local quantum circuits require $\Omega(n)$ depth to be implemented. Our proof gives a matching upper bound for the given IQP circuits.

Recall that IQP gates can be performed in any order (as they commute). Hence, by reordering gates and redefining the $\theta_{i,j}$ angles, any given IQP circuit \mathcal{C} can be put in a normal form \mathcal{C}' that contains at most one single-qubit gate per qubit and one two-qubit gate per pair of qubits. Our approach now is to introduce additional layers of nearest-neighbor SWAP gates following layers of two-qubit gates (Fig. 5). To illustrate the algorithm, we regard qubits as “particles” moving up or down the line by the action of the SWAP gates. At a given step t , we apply a two-qubit IQP gate followed by a SWAP to all pairs of the form $(2i-1, 2i)$ for $1 \leq i \leq \lfloor n/2 \rfloor$ when t is even [resp. $(2i, 2i+1)$ for $1 \leq i \leq \lfloor (n-1)/2 \rfloor$ when t is odd]. By iterating this process n times, the qubit initially in the i th position in the line (with arbitrary i) travels to the $n-i+1$ th position, meeting every other qubit exactly one time along the way because of the intermediate value

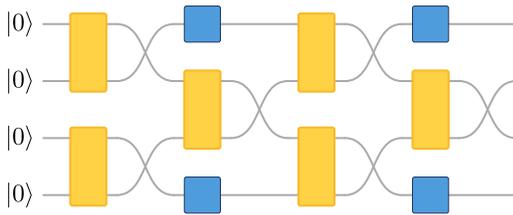


FIG. 5. Linear-depth implementation of dense IQP circuits (10). We illustrate our algorithm for four qubits. Yellow (resp. blue) blocks implement the two- (resp. one-) qubit gates in Eq. (10). The circuit consists of four single-qubit gates (resp. six two-qubit ones), which coincides with the number of vertices (resp. edges) of the complete graph K_4 .

theorem; each two-qubit gate of \mathcal{C}' is implemented in one of these crossings. Furthermore, each qubit spends one step without meeting any qubit when it reaches the line’s boundary; at these points, single-qubit gates can be implemented. \square

2. Proof of Lemma 2

Below, we denote with $\Gamma := \{\beta \in \{0, \theta\}^{\mu mn} : p_\beta \neq 0\}$ the set of allowed configurations for β in step Q1. Let $x \in \{0, 1\}^n$ (resp. $y \in \{0, 1\}^{N-n}$, $N = \mu mn$) be the measurement outcomes of the n rightmost primitive qubits (resp. remaining ones) after step Q3, and let $q(x, y, \beta)$ be the final total probability of observing the values x, y, β .

As in Sec. VIB, it will be convenient to recast our architectures as nonadaptive using mappings (C1) and (C2). In this picture, the following identity readily follows from standard properties of X -teleportation circuits [78,79]:

$$q(x, y, \beta) = q(x, y|\beta)p_\beta = q(x|y, \beta) \frac{1}{2^{N-n}} \frac{1}{|\Gamma|},$$

$$\text{for any } x \in \{0, 1\}^n, y \in \{0, 1\}^{N-n}, \beta \in \Gamma. \quad (12)$$

Above, we used that p_β is uniformly supported over Γ (by design) as well as $q(y|\beta) = 1/2^{N-n}$, which follows from standard properties of X -teleportation circuits [78,79]. Note that $\text{prob}(a, b|\beta)$ in Eq. (3) and $q(x, y|\beta)$ as above are identical probability distributions up to a relabeling $(x, y) = \ell(a, b)$ of the random variables. Thus, if $\tilde{q}(\ell(a, b)|\beta)$ approximates $q(\ell(a, b)|\beta)$ up to relative error $1/4 + o(1)$, then $|\tilde{\mathcal{Z}}^{[ia, (\pi/4)b+\beta]}|^2 := \tilde{q}(\ell(a, b)|\beta)2^{N_x+N_z/2}$ approximates $|\mathcal{Z}^{[ia, (\pi/4)b+\beta]}|^2$ with the same error. Hence, the proof reduces to showing that approximating $q(x, y|\beta)$ for architectures I–III is #P-hard for the given error and $m \in O(n^2)$.

Next, recall that the output probabilities $a_z = |\langle z_1, \dots, z_k | \mathcal{C} | 0 \rangle|^2$ of arbitrary k -qubit dense IQP circuits \mathcal{C} as in Eq. (10) are #P-hard to approximate up to relative error $1/4 + o(1)$ [80,81]. Via Lemmas 3 and 6, the latter can further be implemented in our architectures using lattices with $n \times O(n^2)$ qubits for some $n := k + r$ with $r \in O(k)$; to apply Lemma 3, we can either decompose \mathcal{C} exactly as a 1D homogeneous Clifford + T circuit [82] or use the gadgets in the proofs of Lemmas 4–5 to directly implement the IQP gates. Now, let $|\psi\rangle_{y,\beta}$ denote the state vector of the n rightmost primitive qubits after observing y, β . It follows from our discussion that $|\psi\rangle_{y,\beta} = (\mathcal{C}|0\rangle^{\otimes k})|0\rangle^{\otimes r}$ for some efficiently computable value of y, β . Defining $\bar{z} := (z_1, \dots, z_k, 0_{k+1}, \dots, 0_{k+r})$, it follows that $a_x = q(\bar{z}|y, \beta)$. Furthermore, if $\tilde{q}(\bar{z}, y|\beta)$ approximates $q(\bar{z}, y|\beta)$ up to relative error $\eta > 0$, then $\tilde{q}(\bar{z}|y, \beta) := \tilde{q}(\bar{z}, y|\beta)2^{N-n}$ approximates $q(\bar{z}|y, \beta)$ with the same error. Hence, approximating $q(\bar{z}|y, \beta)$ up to relative error $1/4 + o(1)$ is #P-hard. \square

D. Hardness argument (proof of Theorem 1)

Finally, we prove Theorem 1. Similarly to Refs. [16,17], we apply Stockmeyer's Theorem [73] to relate the problems of approximately sampling from output distributions of quantum circuits to approximating individual output probabilities. Our proof is by contradiction: Assuming that the worst-case #P-hardness of estimating the partition functions of the Ising models (4) extends to the average case (Conjecture 2) and that the output probabilities of architectures I–III are sufficiently anticoncentrated (Conjecture 3), we show that the existence of a classical algorithm for sampling from the latter within a constant ℓ_1 norm implies that a FBPP^{NP} algorithm can solve #P-hard problems; this leads to a collapse of the polynomial hierarchy to its third level, in contradiction with Conjecture 1.

Let $\Gamma := \{\beta \in \{0, \theta\}^{mn} : p_\beta \neq 0\}$ be the set of allowed β configurations in step Q1, let $x \in \{0, 1\}^n$ (resp. $y \in \{0, 1\}^{N-n}$, $N = \mu mn$) be the measurement outcomes of the n rightmost primitive qubits (resp. remaining ones) after step Q3, and let $q(x, y, \beta)$ be the final total probability of observing the values x, y, β . As a preliminary, we prove that if Conjecture 3 holds, then the probability distribution $q(x, y, \beta)$ associated with random input states and measurement outcomes of architectures I–III is anticoncentrated. We first note that $q(x|y, \beta)$ in Eq. (12) coincides with the output distribution of some n -qubit $O(m)$ -depth circuit $\mathcal{C}_{y,\beta}$ of gates of the form (5): This is easily seen from mappings (C1) and (C2) and standard properties of X teleportation [38,78,79] (cf. also the next section and Fig. 6). For arbitrary $y \in \{0, 1\}^{N-n}$, $\beta \in \Gamma$, let us now define

$$\gamma_{y,\beta} := \frac{|\{x \in \{0, 1\}^n : q(x|y, \beta) \geq 1/2^n\}|}{2^n}, \quad (13)$$

$$\alpha := \frac{|\{y \in \{0, 1\}^{N-n}, \beta \in \Gamma : \gamma_{y,\beta} \geq 1/e\}|}{2^{N-n}|\Gamma|}. \quad (14)$$

Here, $\gamma_{y,\beta}$ is the fraction of $\mathcal{C}_{y,\beta}$ output probabilities larger than $1/2^n$, and α is the fraction of $\mathcal{C}_{y,\beta}$ circuits that fulfill Eq. (6). Conjecture 3 states that $\gamma_{y,\beta} \geq 1/e$ for $m \in O(n)$. Consequently, for $n \times O(n)$ lattices, Eq. (12) implies that

$$\text{prob}_{x,y,\beta} \left(q(x, y, \beta) \geq \frac{1}{2^N |\Gamma|} \right) \geq 1/e. \quad (15)$$

Furthermore, since $q(y, \beta)$ is uniformly distributed over its support, it also follows from Eq. (12) that

$$\begin{aligned} \text{prob}_{x,y,\beta} \left(q(x, y, \beta) \geq \frac{1}{2^N |\Gamma|} \right) &= \sum_{y,\beta} \frac{\gamma_{y,\beta}}{2^{N-n} |\Gamma|} \\ &= \mathbb{E}_{y,\beta}(\gamma_{y,\beta}) \geq \frac{\alpha}{e}, \end{aligned} \quad (16)$$

for a uniformly random $x, y \in \{0, 1\}^N$, $\beta \in \Gamma$. Equation (16) tells us that the robustness of the anti-concentration inequality (15) can be tested by computing the average value $\mathbb{E}_{y,\beta}(\gamma_{y,\beta})$ of $\gamma_{y,\beta}$ or by estimating the fraction α . As it is discussed further in Appendix C, $\gamma_{y,\beta} \times e$ and α are expected to converge to 1 for universal 1D nearest-neighbor circuits as n grows asymptotically [43–45,83] in the regime $m \in O(n)$ [84–86]. In Appendix C, Fig. 7, we present numerical evidence that $\mathbb{E}_{y,\beta}(\gamma_{y,\beta}) \rightarrow 1/e$, $\alpha \rightarrow 1$ in the asymptotic limit and a tight agreement for $n \geq 9$; more strongly, we also find that $\gamma_{y,\beta}$ is nearly $1/e$ for almost every uniformly sampled instance for $n \geq 9$ and that $q(x|y, \beta)$ is Porter-Thomas distributed, which is a signature of Haar-random chaotic unitary processes [18,41–45].

We are now ready to prove Theorem 1. For any architecture I–III, we let a denote an element of $\{0, 1\}^N \times \Gamma$ and assume that the output distribution $p_c(a)$ of a classical BPP algorithm fulfills

$$\sum_{a \in \{0, 1\}^N \times \Gamma} |p_c(a) - q(a)| \leq \varepsilon, \quad (17)$$

for a constant $\varepsilon \geq 0$. By Stockmeyer's Theorem [73] and the triangle inequality, there exists a FBPP^{NP} algorithm that computes an estimate $\widetilde{p}_c(a)$ such that

$$\begin{aligned} &| \widetilde{p}_c(a) - q(a) | \\ &\leq \frac{q(a)}{\text{poly}(N)} + |p_c(a) - q(a)| \left(1 + \frac{1}{\text{poly}(N)} \right), \end{aligned}$$

where we used $\log |\Gamma| \in O(N)$ to remove dependencies on $|\Gamma|$. From Markov's inequality and Eq. (17), we get that

$$\text{prob}_a \left(|p_c(a) - q(a)| \geq \frac{\varepsilon}{2^N |\Gamma| \delta} \right) \leq \delta, \quad (18)$$

for any constant $0 < \delta < 1$, where $a \in \{0, 1\}^N \times \Gamma$ is picked uniformly at random. Hence,

$$| \widetilde{p}_c(a) - q(a) | \leq \frac{q(a)}{\text{poly}(N)} + \frac{\varepsilon(1 + o(1))}{\delta 2^N |\Gamma|} \quad (19)$$

with probability of at least $1 - \delta$ over the choice of a .

We now claim that Eqs. (19) and (15) simultaneously hold for a single $q(a)$ with probability $(1/e)(1 - \delta)$ since a classical description of I–III does not reveal to a classical sampler which output probabilities are #P-hard to approximate: Hence, the latter cannot adversarially corrupt the latter. This is manifestly seen at the encoded random circuit level because of the presence of random by-product

operators of the form $\prod_i X_i^{y_i}$ (with random y_i), which obfuscate the location of the #P-hard probabilities from the sampler [17,47]. Hence, setting $\varepsilon = \gamma/8, \delta = \gamma/2, \gamma = 1/e$, we obtain

$$|\widetilde{p_c}(a) - q(a)| \leq \left(\frac{1}{4} + o(1)\right)q(a), \quad (20)$$

with probability of at least $\gamma(1 - \gamma/2) > 0.3$ over the choice of a . Setting $\varepsilon = 1/22 < \gamma/8$, the above procedure yields an approximation $\widetilde{p_c}(a)$ of $q(a)$ up to relative error $1/4 + o(1)$. Using Eqs. (3) and (12), we obtain a FBPP^{NP} algorithm that approximates $|\mathcal{Z}^{\alpha,\beta}|^2$ with relative error $1/4 + o(1)$ for at least a 0.3 fraction of the instances. This yields a contradiction.

As final remarks, note that the above argument is robust to small finite-size variations to the threshold $\gamma = 1/e$ in Conjecture 3, Eq. (6), since the constants $\varepsilon = \gamma/5, \delta = \gamma/2$, and $\gamma(1 - \gamma/2)$ have only linear and quadratic dependencies on γ . Also, notice that Conjectures 2 and 3 enter the above argument in order to allow for a constant additive error ε , which is key for a real-life demonstration of a quantum speedup. Additive errors give rise to demanding, but not unrealistic prescriptions. However, in the ideal case where one assumes no sampling errors, or multiplicative sampling errors, our result holds even without these conjectures via the arguments in Refs. [21,22].

VII. CONCLUSION

In this work, we have established feasible and simple schemes for quantum simulation that exhibit a superpolynomial quantum speedup with high evidence, in a complexity-theoretic sense. As such, this work is expected to significantly contribute to bringing notions of quantum devices outperforming classical supercomputers closer to reality. This work can be seen as an invitation towards a number of further exciting research directions: While the schemes presented may not quite yet constitute experimentally realizable blueprints, it should be clear that steps already experimentally taken are very similar to those discussed. Hence, it seems interesting to explore detailed settings for cold atoms or trapped ions in detail, requiring little local control and allowing for comparably short coherence times. What is more, it appears obvious that further complexity-theoretic results on intermediate problems are needed to fully capture the potential of quantum devices outperforming classical computers without being universal quantum computers. It is our hope that the present work can contribute to motivating such further work, guiding experiments in the near future.

ACKNOWLEDGMENTS

We thank Scott Aaronson, Dan E. Browne, Andreas Elben, Bill Fefferman, Wolfgang Lechner, and Ashley Montanaro for discussions. J.B.V. thanks Vadym Kliuchnikov and Neil Julien Ross for helpful comments. This work was supported by the EU Horizon 2020 (640800 AQuS), the Templeton and Alexander von Humboldt Foundations, the ERC (TAQ), and the DFG (CRC 183, EI 519/7-1). R. R. is funded by NSERC, Cifar, IARPA.

APPENDIX A: FULL HAMILTONIAN (2)

For any architecture I–III (Fig. 1), let $\mathcal{L}_P = (V_P, E_P)$ be the sublattice of \mathcal{L} containing all primitive qubits (i.e., the square lattice subgraph of \mathcal{L}). Furthermore, let $\mathcal{L}_{DB} = (V_{DB}, E_{DB})$ be \mathcal{L} 's sublattice containing all dangling bonds for architecture III and the empty graph otherwise. For any $i \in V$, let $\text{deg}_P(i)$ [resp. $\text{deg}_{DB}(i)$] be the number of primitive (resp. dangling-bond) qubits connected to i in \mathcal{L} . Then, the full Hamiltonian (2) of the experiment reads

$$H = \underbrace{\sum_{(i,j) \in E_P} \frac{\pi}{4} Z_i Z_j - \sum_{i' \in V_P} \frac{\pi}{4} \text{deg}_P(i') Z_{i'}}_{H_{CZ}} + \underbrace{\sum_{(k,l) \in E_{DB}} \frac{\pi}{16} Z_k Z_l - \sum_{k' \in V_{DB}} \frac{\pi}{16} \text{deg}_{DB}(k') Z_{k'}}_{H_{CT}}. \quad (\text{A1})$$

Above, the Hamiltonian H_{CZ} (resp. H_{CT}) implements a CZ (resp. CT) gate on every edge of the bright (resp. dark) sublattice in Fig. 1. Note that H_{CT} is not present in architectures I and II. Also, realize that $\text{deg}_P(i)$ takes value 4 on bulk qubits, 2 on the corners, and 3 elsewhere on edges; for architecture III, $\text{deg}_{DB}(i)$ takes value 1 everywhere.

APPENDIX B: MAPPING OUTPUT PROBABILITIES TO ISING PARTITION FUNCTIONS

Let $\mathcal{L}_X = (V_X, E_X)$ [resp. $\mathcal{L}_Z = (V_Z, \emptyset)$] be the primitive-qubit square sublattice of \mathcal{L} (resp. the disjoint union of all dangling-bond qubits), and pick \mathcal{L}_X to be the lattice \mathcal{L}_{sq} in Sec. III, Eq. (4). Let $\alpha := \pi a, \vartheta := \beta + (\pi/4)b$, where we let b be the string of outcomes of the Z measurements in architecture III, and define $b_i := 0, i = 1, \dots, N_Z$, by convention, for architectures I and II. Furthermore, let $\theta = \pi/8$ for architecture II and $\theta = \pi/4$ otherwise. We now prove Eq. (3) using formula (A1):

$$\begin{aligned}
|\langle a, b | e^{-iH} | \psi_\beta \rangle| &= \left| \langle a, b | e^{-i[H + \sum_{i \in V} (\beta_i/2) Z_i]} | + \rangle^{\otimes N} \right| = \left| \langle a, b | e^{-i\{\sum_{(i,j) \in E} J_{ij} Z_i Z_j - \sum_{i \in V} [h_i - (\beta_i/2)] Z_i\}} | + \rangle^{\otimes N} \right| \\
&= \frac{1}{\sqrt{2^{N_Z}}} \left| \langle a | e^{-i\{\sum_{(i,j) \in E_X} (\pi/4) Z_i Z_j - \sum_{i \in V_X} \sum_{(i,k) \in E_Z} [h_i - (\beta_i/2) - [(\pi b_i)]/8] Z_i\}} | + \rangle^{\otimes N_X} \right| \\
&= \frac{1}{\sqrt{2^{N_Z}}} \left| \langle + | e^{-i\sum_{i \in V_X} (\alpha_i/2) Z_i} e^{-i\{\sum_{(i,j) \in E_X} (\pi/4) Z_i Z_j - \sum_{i \in V_X} \sum_{(i,k) \in E_Z} [h_i - (\vartheta_i/2)] Z_i\}} | + \rangle^{\otimes N_X} \right| \\
&= \frac{1}{\sqrt{2^{N_Z}}} \left| (\langle 0 | H)^{\otimes N_X} e^{-iH^{(\alpha,\beta)}} (H | 0 \rangle)^{\otimes N_X} \right| = \frac{1}{\sqrt{2^{N_Z} 2^{N_X}}} \left| \sum_{x,y \in \{0,1\}^{N_X}} \langle x | e^{-iH^{(\alpha,\beta)}} | y \rangle \right| \\
&= \left| \frac{\text{tr}(e^{-iH^{(\alpha,\beta)}})}{\sqrt{2^{N_Z+2N_X}}} \right| = \frac{|\mathcal{Z}(\alpha,\beta)|}{\sqrt{2^{N_Z+2N_X}}},
\end{aligned}$$

where we have defined $|+\rangle := (|0\rangle + |1\rangle)/\sqrt{2}$ in the second step and used that $H^{(\alpha,\beta)}$ is diagonal in the final one. We obtain Eq. (3) by squaring; therein, $\alpha_i \in \{0, \pi\}$, $\vartheta \in \{0, \theta\}$ follows from the definition.

APPENDIX C: NUMERICAL EVIDENCE FOR ANTICENTRATION OF THE OUTPUT DISTRIBUTION

In this appendix, we present numerical evidence for the validity of Conjecture 3, which is exploited in the proof of Theorem 1 as discussed in Sec. VID. Therein, we discussed that for any architecture I–III, given an initial n -row, m -column square lattice, there exists an n -qubit D -depth circuit family $\{\mathcal{C}_{\beta,y}\}_{y,\beta}$ of gates of the form (5), with $D \in O(m)$, such that $q(x|y, \beta) = |\langle x | \mathcal{C}_{\beta,y} | 0 \rangle|^2$. If the circuits $\{\mathcal{C}_{\beta,y}\}_{y,\beta}$ exhibit anticentration as in Conjecture 3, then $q(x, y, \beta)$ is anticentrated as in Eq. (15). This is used in Sec. VID to turn an approximate classical sampler into a FBPP^{NP} algorithm to approximate single output probabilities with high accuracy.

The concrete circuit families associated with each architecture are derived below and depicted in Fig. 6. To numerically test Conjecture 3, we have performed simulations of randomly generated circuits of gates of the form

(5) (for each circuit family) in LIQUiD [87] with up to 20 logical qubits. For each system size, we generated 100 random instances for circuits associated with $n \times n$ and $n \times n^2$ lattices. For each instance, we evaluated exactly the fraction $\gamma_{y,\beta}$ [Eq. (13)] of output probabilities fulfilling Eq. (6). Our results are summarized in Fig. 7: Therein, one can see that for circuits associated with both $n \times n$ and $n \times n^2$ lattices, this fraction quickly approaches a constant $\gamma = 1/e$ with rapidly decreasing variance with respect to the choice of circuits. We can conclude that, with very high probability, in a realization of the proposed experiment, the amplitude of the final state of the computation anticentrates.

As discussed in Refs. [17,20,47], it might seem *a priori* counterintuitive that constant-depth nearest-neighbor architectures anticentrate. However, the above connections between our architectures and random circuits shed key insights into why this behavior is actually natural. As shown in Sec. VIB, the random logical circuits of gates of the form (5) encoded in our architectures are universal for quantum computation. Universal random quantum circuits of increasing depth are known to approximate the Haar measure under various settings [43–45,83,86]. For 1D nearest-neighbor layouts, the latter are expected to reach a chaotic Porter-Thomas-distributed regime [41–43] in

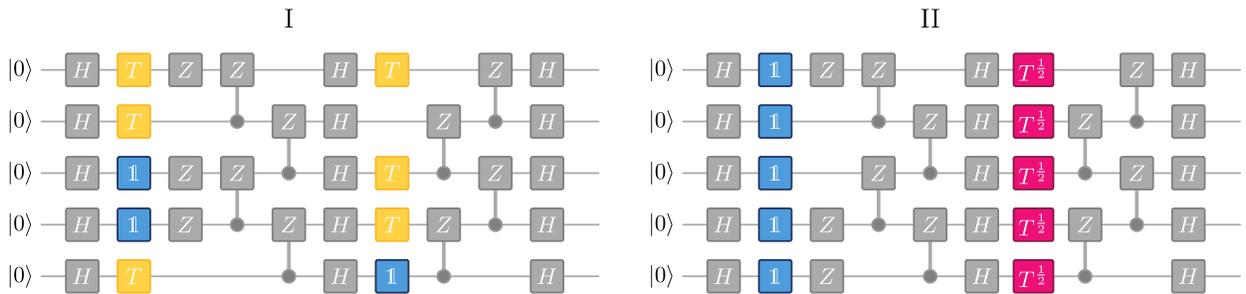


FIG. 6. We show the logical circuit corresponding to architectures I and III (left) and II (right) for 5-row, 2-column lattices prepared with uniformly random $\beta_i \in \{0, \pi/4\}$ (blue, yellow) (I and III) and column-random $\beta_i \in \{0, \pi/8\}$ (blue, magenta) (II). At those qubits that have been prepared with $\beta_i = 0$, an identity gate is applied; on those with $\beta_i = \pi/4$, a T gate; and on those with $\beta_i = \pi/8$, a \sqrt{T} gate.

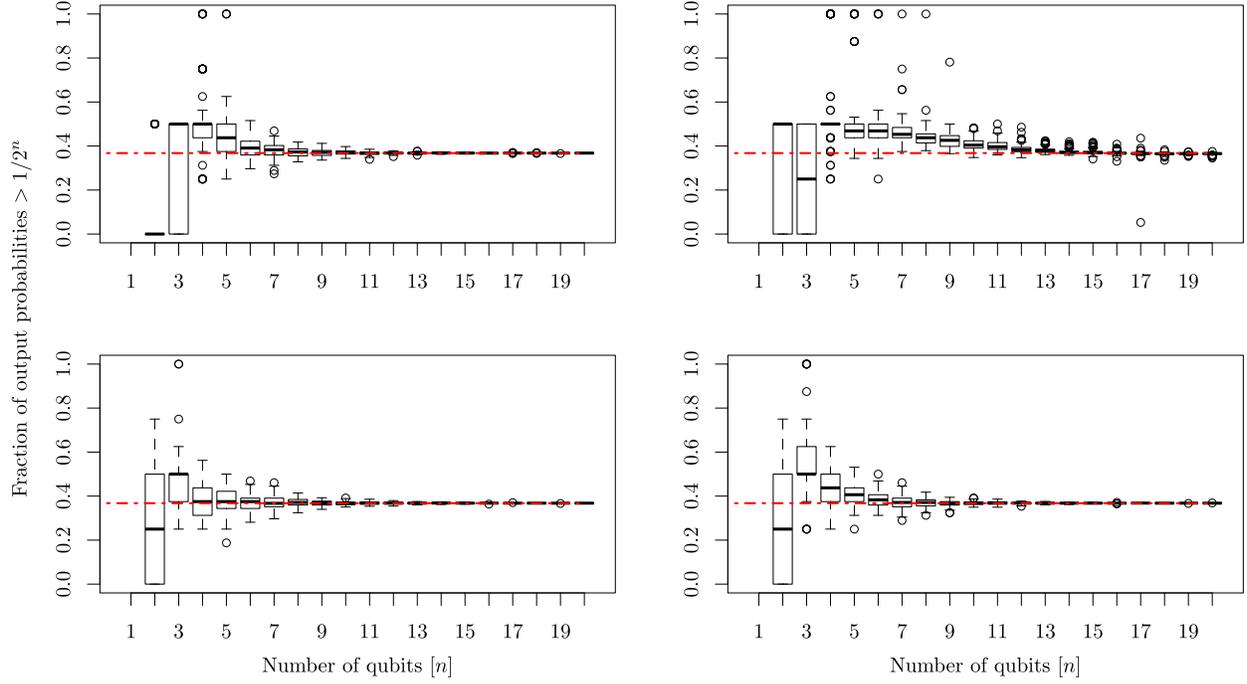


FIG. 7. Fraction of output probabilities larger than $1/2^n$ of random circuits drawn from the families \mathcal{F}_{DO} (l.h.s.) and \mathcal{F}_{col} (r.h.s.) for both linear (top) and quadratic (bottom) circuit depth in the number of qubits the circuit acts upon n , i.e., lattices of size $n \times n$ (top) and $n \times n^2$ (bottom). For each n , we draw 100 i.i.d. realizations (β, y) , thus of the circuit $\mathcal{C}_{\beta,y}$, and plot the resulting distribution in the form of a box plot. The red dashed line shows the value of $1/e$, which is precisely the value to be expected if the output probabilities are Porter-Thomas distributed.

depth $D \in O(n)$ [84–86] (cf. Ref. [18] for further discussion). As an additional piece of supporting evidence for anticoncentration, we numerically confirmed that our output probabilities are close to being Porter-Thomas distributed in the ℓ_1 norm (Fig. 8). Furthermore, our numerics are in agreement with prior numerical works on MBQC settings [45] and other gate sets in 2D layouts [18,46].

Circuit families.—For the sake of completeness, we spell out the logical circuits that are effectively implemented in architectures I and II. The latter are derived via mappings (C1) and (C2) and X -teleportation properties [78,79]. Examples for 4×2 lattices are depicted in Fig. 6. The logical circuit family corresponding to I and III (resp. to II) is denoted \mathcal{F}_{DO} (resp. \mathcal{F}_{col}). Let us now label primitive-lattice sites by row-column coordinates $[i, j]$. The circuits are generated inductively, starting from the left column $j = 1$. Measurements are ordered from left to right. The computation begins on the $|+\rangle^{\otimes n}$ state and proceeds as follows:

- (1) Apply the gate $\exp(i\beta_{[i,j]}Z_{[i,j]})$ to qubit $[i, j]$, with $\beta_{[i,j]}$ chosen as in step I for II and III; for III, we let $\beta_{[i,j]} := s_{[i,j]}\pi/4$, where $s_{[i,j]}$ is the outcome after measuring the dangling neighbor of $[i, j]$.
- (2) If $j < m$, apply a random $Z_{[i,j]}^{a_{[i,j]}}$ gate to every qubit $[i, j]$, where $a_{[i,j]}$ is the outcome of the measurement at site $[i, j]$.
- (3) Apply CZ on all neighboring qubits.
- (4) Apply a Hadamard gate to each qubit.

- (5) If $j = m$, measure in the standard basis and terminate; otherwise, increase $j := j + 1$.

a. Convergence to the chaotic regime. It is an interesting detail that the value of $\gamma = 1/e$, which we observe above, is a signature of the exponential distribution (also known as Porter-Thomas distribution) that is known to emerge in chaotic quantum systems for large system sizes [18,41–43, 45,46]. This distribution is given by

$$P_{PT}(p) = 2^n \exp(-2^n p) \tag{C1}$$

and thus anticoncentrates in precisely the fashion observed here. Note that the same behavior was observed in previous work investigating random MBQC settings [45], as well as recently in Ref. [18], which investigated random universal circuits on a 2D architecture. Notably, the finite and universal gate sets considered in these works are very similar to the ones considered here. Likewise, convergence to the exponential distribution was observed in Ref. [46] for approximately Haar-random two-qubit unitaries in a 2D setup.

In Fig. 8, we show the total variation distance between the empirical distributions of output probabilities of the random circuits generated in our numerical experiments and the discretized Porter-Thomas distribution. We can see that as the number of qubits increases, the output distributions of random circuits approach Porter-Thomas distribution.

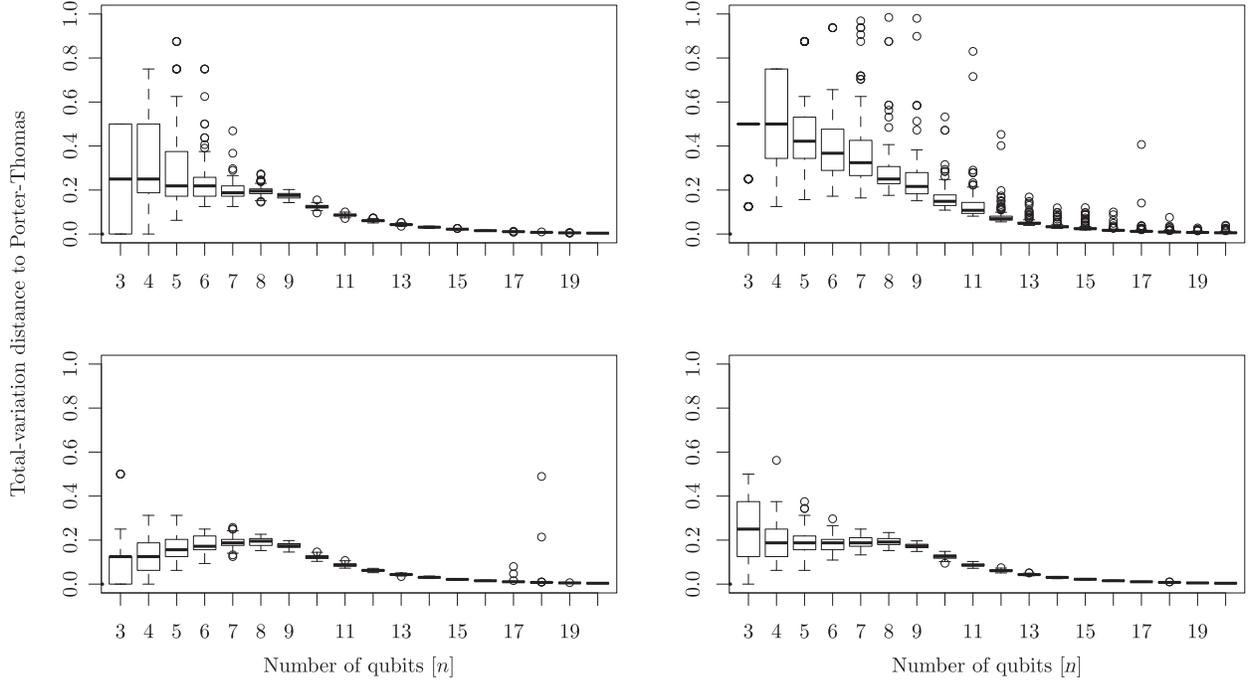


FIG. 8. Total variation distance to the Porter-Thomas distribution of the empirical distribution of output probabilities of random circuits from the families \mathcal{F}_{DO} (l.h.s.) and \mathcal{F}_{col} (r.h.s.) for both linear (top) and quadratic (bottom) circuit depths in the number of qubits the circuit acts on n , i.e., lattices of size $n \times n$ (top) and $n \times n^2$ (bottom). For each n , we draw 100 i.i.d. realizations (β, y) , thus of the circuit $\mathcal{C}_{\beta, y}$, and plot the resulting distribution in the form of a box plot.

To calculate the total variation distance to the exponential distribution (C1), we discretized the interval $[0, 1]$ into m bins, each of which contains probability weight $1/m$. In other words, the discretization (p_0, p_1, \dots, p_m) is defined by given $p_0 = 0$, $p_m = 1$ and

$$\int_{p_i}^{p_{i+1}} P_{PT}(p) dp = \frac{1}{m}. \quad (\text{C2})$$

Denote by $Q(p)$ the numerically observed distribution of output probabilities $p = |\langle x | \mathcal{C} | 0 \rangle|^2$ over the set $\Omega = \{[p_i, p_{i+1}]\}_{i=0, \dots, m}$. The total variation distance between P and the exponential distribution is then given by

$$\|P - Q\|_{\text{TV}} = \frac{1}{2} \sum_{X \in \Omega} |P(X) - 1/m|. \quad (\text{C3})$$

Since the number of samples we obtain in each run is given by 2^n , we choose the number of bins m depending on n . Specifically, we choose $m = \min\{\lceil 2^n/5 \rceil, 100\}$ to allow fair comparison for small n .

APPENDIX D: #P-HARDNESS FROM $n \times O(n)$ LATTICES

In this appendix, we show that #P-hardness of approximating output probabilities in Lemma 2 arises already for $n \times m$ -qubit lattices with $m \in O(n)$. Specifically, we prove this by introducing two slight modifications in architectures

I and II (cf. Fig. 4) without changing the fundamental structure of the basic layout of steps Q1–Q3.

The key idea is to introduce different types of input states (with different β_i 's) in the preparation step Q3: Specifically, we pick $\beta_i \in \{0, \theta\}$ for qubits on odd-row sites, $\beta_i \in \{0, \pi/2\}$ on even-row ones; additionally, we perform a local Hadamard rotation on even-row sites before the Ising Hamiltonian evolution in step Q2 begins. This is shown in Fig. 9. The net effect is to initialize even-site qubits on either $|0\rangle$ (the $+1$ eigenstate of Z) or $|-i\rangle := |0\rangle - i|1\rangle$ (the



FIG. 9. Modified architectures I and II, named “MI” and “MII” in the figure. Changes are introduced on the even rows with respect to Fig. 1. On even rows, we pick $\beta_i \in \{0, \pi/2\}$ (black stands for 0, white stands for $\pi/2$) and perform a local Hadamard rotation. In architecture MI, β_i is uniformly random. In architecture MII, β_i is translation invariant on columns with a period less than or equal to 4 (i.e., $\text{TI}_{(4, \infty)}$ in the notation of the main text).

-1 eigenstate of Y) at random. Qubits initialized in $|0\rangle$ are invisible to the Ising evolution Q2 and effectively become unentangled from the computation. Furthermore, preparing a qubit in $| - i \rangle$, evolving Q2 and measuring X is equivalent to preparing $| + \rangle$ and measuring Y instead at the end of the computation. Again, we choose the $|\psi_\beta\rangle$ to be fully disordered (DO) for architecture I. For architecture II, we pick $|\psi_\beta\rangle$ to be TI in one direction with a period of at most 4, i.e., $\text{TI}_{(4,\infty)}$ symmetric in our notation.

The full experiment can now be mapped to the non-adaptive MBQC analogue to Eq. (C1) with two differences: First, the MBQC acts on a graph-state vector $|G\rangle$ [88], instead of a cluster state, whose underlying graph G is derived from the 2D lattice by deleting $|0\rangle$ -state vertices (the output probabilities of the computation can be mapped to an Ising model on G using the tools of Appendixes A and B); second, the remaining vertices on even columns are measured on the Y basis. We now pick $n = 2k - 1$ and study the logical circuit of the MBQC in two scenarios:

- (i) All even-row qubits are initialized in $|0\rangle$. The MBQC acts on a graph-state vector $|G'\rangle$ that is the product of k disconnected 1D cluster states. Modulo by-product operators, the local measurements drive a random logical k -qubit circuit of single-qubit gates $\{R_i^{a_i}(\theta) := H_i e^{i(a_i\theta)Z_i}, a_i \in \{0, 1\}\}$ and depth $m - 1$ (information flows on odd rows). For architecture II, the latter circuit inherits a $\text{TI}_{(2,\infty)}$ symmetry from the $\text{TI}_{(4,\infty)}$ one of the input-state vector $|\psi_\beta\rangle$.
- (ii) Even-column, even-row qubits are initialized in $|0\rangle$; even-column, odd-row ones are left unspecified. The MBQC acts on a cluster state with ‘‘holes’’ $|G''\rangle$ as in Refs. [89,90]. Information flows again on odd rows. If an even-column, odd-row qubit i is prepared in $| - i \rangle$ and measured in the X basis (or, equivalently, in $| + \rangle$ and measured in the Y basis), we obtain a

reduced graph-state vector $|G'''\rangle$ whose graph G''' is obtained from G'' by contracting the edges incident to i [89,90]. Thus, $| - i \rangle$ state vectors between the odd qubit lines let us implement logical entangling gates of the form

$$E^{b,c}(\theta) := \left(\prod_{i=1}^k H_i \right) \left(\prod_{i=1}^{k-1} CZ_{i,i+1}^{b_i} \right) \left(\prod_{i=1}^k e^{-i(c_i\theta)Z_i} \right),$$

$$b_i, c_i \in \{0, 1\},$$

where $b_i = 1$ if the qubit between lines $2i, (2i - 1)$ is in $| - i \rangle$ and zero otherwise; and c_i indicates whether we measure X or $X_{-\theta}$ on the $(2i - 1)$ th line. Postselecting b_i gives us the ability to implement nontranslation-invariant two-qubit entangling gates between qubit lines at will. Again, for architecture II, the gate $E^{b,c}(\theta)$ inherits a $\text{TI}_{(2,\infty)}$ symmetry.

Combining the above facts, it follows that we can simulate arbitrary k -qubit, nearest-neighbor, Clifford + T circuits in the modified architecture I via postselection, using lattices with $(2k - 1) \times (2k - 1)$ qubits. The latter can efficiently implement the #P-hard IQP circuits of Lemma 6 with a constant-overhead factor.

In the modified architecture II, observation (i) and postselection of by-product operators (Fig. 4) yields $\text{TI}_{(2,\infty)}$ -symmetric circuits of $\{H_i, e^{\pm i(\pi/16)Z_i}, e^{\pm i(\pi/16)X_i}\}$ gates. In combination with the gadgets in the proof of Lemma 4, this lets us implement k -qubit, $\text{TI}_{(2,\infty)}$ -symmetric, nearest-neighbor circuits of $e^{\mp i(\pi/8)X_i X_{i+1}}$ gates. Furthermore, by-product operators also let us break the $\text{TI}_{(2,\infty)}$ symmetry via the identities below and allow us to implement non-TI arbitrary, nearest-neighbor, Clifford + T circuits as in the previous case with constant overhead:

$$\prod_{i=1}^{\lfloor k/2 \rfloor} e^{ia_i(\pi/4)X_{[i]}X_{[i+1]}} = \left(\prod_{i=1}^{\lfloor k/2 \rfloor} e^{i(\pi/8)X_{[i]}X_{[i+1]}} \right) \left(\prod_{i=1}^{\lfloor k/2 \rfloor} Z_{[i]}^{a_i} \right) \left(\prod_{i=1}^{\lfloor k/2 \rfloor} e^{-i(\pi/8)X_{[i]}X_{[i+1]}} \right), \quad [i] := 2i - x_0, x_0 \in \{0, 1\}, a_i \in \{0, 1\},$$

$$\prod_{i=1}^{\lfloor k/2 \rfloor} e^{ib_i(\pi/8)X_{[i]}} = \left(\prod_{i=1}^{\lfloor k/2 \rfloor} e^{i(\pi/16)X_{[i]}} \right) \left(\prod_{i=1}^{\lfloor k/2 \rfloor} Z_{[i]}^{b_i} \right) \left(\prod_{i=1}^{\lfloor k/2 \rfloor} e^{-i(\pi/16)X_{[i]}} \right), \quad [i] := 2i - x_0, x_0 \in \{0, 1\}, b_i \in \{0, 1\}.$$

Again, this yields an efficient and exact postselected implementation of the desired IQP circuits on k logical qubits.

APPENDIX E: CERTIFICATION PROTOCOL

In this appendix, we describe an efficient parallelizable certification protocol for ground states of gapped local Hamiltonians. The protocol is an optimization of the algorithm presented in Ref. [51] featuring a reduced sample complexity. A direct implementation of the algorithm of

Ref. [51] requires a supercubic number $O(N^3 \log N)$ of prepare-and-measure experiments, where N is the system size. Our protocol brings this complexity down to $O(N^2)$ by combining parallel sequences of local measurements and efficient classical postprocessing. Additionally, it relies on on-site single-qubit observables only, as opposed to few-body Hamiltonian terms, which are needed in Ref. [51]. These improvements render the protocol of Ref. [51] faster and more suitable for, e.g., near-term quantum speedup experiments. However, our optimized algorithm can be

used for verifying ground states of arbitrary nondegenerate gapped local Hamiltonians and may also be of independent interest.

1. Parent Hamiltonians

As a preliminary, we derive nondegenerate gapped local parent Hamiltonians for the premeasurement state vectors $|\psi_\beta\rangle$ of our architectures. To find the stabilizer operators corresponding to $|\psi_\beta\rangle$ in architectures I–III, we note that $|\psi_\beta\rangle$ can be prepared by applying the tensor-product unitary $U_\beta = \prod_{i \in V} e^{-\beta_i Z_i / 2}$ followed by the local quench unitary U to the initial state $|+\rangle_V = \prod_{i \in V} |+\rangle_i$, where $|+\rangle = (|0\rangle + |1\rangle) / \sqrt{2}$ as $|\Psi_\beta\rangle = UU_\beta|+\rangle_V$. Since $|+\rangle$ is a +1 eigenvector of X , the parent Hamiltonian of $|\Psi_\beta\rangle$ is given by $-UU_\beta(\sum_{i \in V} X_i)U_\beta^\dagger U^\dagger$. For architectures I and II, this evaluates to

$$H_{\text{I,II}} = -\sum_{i \in V} \left(X_{\beta_i, i} \prod_{j: (i,j) \in E} Z_j \right). \quad (\text{E1})$$

For the specific case of architecture III, let us partition the vertices of the lattice \mathcal{L} into two sets V_1 and V_2 , such that V_1 contains all qubits on the square primitive sublattice and V_2 the remaining dangling-bond qubits. Furthermore, let E_1 contain those edges where $J_{i,j} = \pi/4$, and E_2 the dangling bonds where $J_{i,j} = \pi/16$. We find the corresponding parent Hamiltonian to be

$$H = -\sum_{\substack{i \in V_1 \\ k: (i,k) \in E_2}} \left(CT_{(i,k)} X_i CT_{(i,k)}^\dagger \prod_{j: (i,j) \in E_1} Z_j \right) - \sum_{\substack{k \in V_2 \\ i: (i,k) \in E_2}} (CT_{(i,k)} X_k CT_{(i,k)}^\dagger), \quad (\text{E2})$$

where the two-body terms evaluate to

$$CT_{(i,k)} X_i CT_{(i,k)}^\dagger = (X_i - TX_i T^\dagger) Z_k + (X_i + TX_i T^\dagger) I_k. \quad (\text{E3})$$

The Hamiltonian terms of H_{III} are 6-local except at the boundary where its locality is reduced.

2. Energy estimation of G -local Hamiltonians

An N -qubit nondegenerate gapped local Hamiltonian $H = \sum_{i \in V} h_i$ with τ -body interactions on a (simple connected) interaction graph $G = (V, E)$ is called “ G -local” if each qubit is located at a vertex $i \in V$ and each term h_i is supported on the neighborhood $\partial(i)$ of i ,

$$\text{supp}(h_i) \subseteq \partial(i) := \{j: (i,j) \in E\}, \quad \forall i \in V. \quad (\text{E4})$$

Though we consider arbitrary interaction graphs G in our analysis, we are particularly interested in constant-degree

ones in our applications: i.e., those with maximum vertex degree $\deg(G)$ upper bounded by a constant, independently of the number of qubits. Because, w.l.o.g., we can pick $\tau = \deg(G)$, the latter graphs model physical systems with geometrically constrained connectivity, which are ubiquitous both in condensed matter physics and quantum information processing. Examples of such graphs are lattices of constant geometric dimension D and fixed-size primitive cells. In particular, the Hamiltonians of the main text are \mathcal{L} local and have $\deg(\mathcal{L}) = 4$ ($\deg(\mathcal{L}) = 5$) for architectures I and II (III).

The key ingredient of our result below is a subroutine for estimating average energies of G -local Hamiltonians using parallelized measurement circuits with time complexity dominated by chromatic number $G^{(2)} = (V^{(2)}, E^{(2)})$ of the next-neighbor interaction graph of $G^{(2)}$. The latter has with same vertices $V^{(2)} = V$ as G and edges between all pairs of neighbors and next neighbors of G [$(v_1, v_2) \in E^{(2)}$ iff $v_1, v_2 \in V$ and have graph distance $d(v_1, v_2) \leq 2$]. The existence of such parallel circuits relies on the existence of certain decompositions of local Hamiltonians into commuting terms, named (κ, α, τ) (“cat”) decompositions below. In short, a Hamiltonian is (κ, α, τ) if it can be decomposed as small sum Hermitian operators that admit parallel measurement circuits that use single-shot τ -body operators.

Definition 7 ((κ, α, τ) decomposition) A G -local N -qubit Hamiltonian H on a graph $G = (V, E)$ is (κ, α, τ) decomposable if there exist constants $\alpha \in (0, 1]$, $\kappa, \tau, \chi \in \mathbb{N}$, a partition $V = \bigcup_{i=1}^{\kappa} V_i$, $\chi \leq \kappa$, and a map $f: [\kappa] \rightarrow [\chi]$, such that

$$H = \sum_{i=1}^{\kappa} H_i, \quad H_i := \sum_{j \in V_{f(i)}} h_j^{(i)}, \quad \text{supp}(h_j^{(i)}) \subset \partial(j), \quad (\text{E5})$$

where (a) $\max_i |V_i| \leq \alpha N$, (b) the terms $h_j^{(i)}$ are Hermitian, and (c) the energy distribution of H_i can be sampled from parallel measurements of τ -body observables and efficient classical postprocessing.

A sufficient condition for Definition 7 (c) to hold is that $i \in [\kappa]$ and that all terms in $\{h_j^{(i)}\}_j$ are τ local and have nonoverlapping support. Then, one can simply measure all $h_j^{(i)}$ in parallel, obtain the outcomes $\{e_j^{(i)}\}_j$, and sample the output distribution of H_i by classically computing the sum $e_i := \sum_{j \in V_{f(i)}} e_j^{(i)}$. The fact that Hamiltonians of this form admit parallel measurement circuits is formalized by the following lemma, which generalizes Lemma 1 in Ref. [51].

Lemma 8 (Estimation of the energy) Let H be an N -qubit Hamiltonian with a given (κ, α, τ) decomposition (E5). Let $P_{i,\mu}$ be the μ th $e_{i,\mu}$ eigenprojector of H_i , for $i \in [\kappa]$, and $X_i^{(j)}$ be the random variable that takes the value $e_{i,\mu}$ with probability $\text{tr}(\rho_p^{(j)} P_{i,\mu})$, modeling a

measurement of H_i on the j th copy of ρ_p . Moreover, let $\langle H_i \rangle_{\rho_p}^* = (1/m) \sum_{i=1}^m X_i^{(j)}$ be the estimate of $\langle H_i \rangle$ on ρ_p by a finite-sample average of m measurement outcomes, and $\langle H \rangle_{\rho_p}^* = \sum_{i=1}^{\kappa} \langle H_i \rangle_{\rho_p}^*$ the resulting estimate of $\langle H \rangle_{\rho_p}$. Last, let $J = \max_{\lambda} \|h_{\lambda}\|$. Then, for any $\overline{p_{\text{err}}} \in [1/2, 1)$ and $\epsilon > 0$, it holds that

$$\mathbb{P}[\|\langle H \rangle_{\rho_p}^* - \langle H \rangle_{\rho_p}\| \leq \epsilon] \geq \overline{p_{\text{err}}}, \quad \text{whenever} \\ m \geq \left[\frac{\alpha^2 \kappa^2 J^2}{2\epsilon^2} \ln \left[\frac{\kappa + 1}{\ln(1/\overline{p_{\text{err}}})} \right] \right] N^2. \quad (\text{E6})$$

Proof.—Since the random variables $\{X_i^{(j)}\}_j$ are independent and $0 \leq X_i^{(j)} \leq \|H_i\|$, Hoeffding's inequality implies

$$\forall i \in [\kappa]: \mathbb{P}[\|\langle H_i \rangle_{\rho_p}^* - \langle H_i \rangle_{\rho_p}\| \leq \epsilon] \geq 1 - 2e^{-[(2m\epsilon^2)/(\|H_i\|^2)]},$$

and since all measurements are independent,

$$\mathbb{P}[\|\langle H \rangle_{\rho_p}^* - \langle H \rangle_{\rho_p}\| \leq \epsilon] \\ \geq \mathbb{P}\left[\forall i \in [\kappa]: \|\langle H_i \rangle_{\rho_p}^* - \langle H_i \rangle_{\rho_p}\| \leq \frac{\epsilon}{\kappa}\right] \\ \geq \prod_{i=1}^{\kappa} (1 - 2e^{-[(2m\epsilon^2)/(\kappa^2 \|H_i\|^2)]}) \\ \geq (1 - 2e^{-[(2m\epsilon^2)/\kappa^2 (\alpha N J)^2]})^{\kappa} \geq \overline{p_{\text{err}}}.$$

The latter identity holds whenever $m \geq m_{\text{opt}}$, with

$$m_{\text{opt}} := \frac{\alpha^2 \kappa^2 J^2}{2\epsilon^2} \ln \left[\frac{2}{1 - \overline{p_{\text{err}}}^{1/\kappa}} \right] N^2 \\ \stackrel{[91]}{\leq} \left[\frac{\alpha^2 \kappa^2 J^2}{2\epsilon^2} \ln \left[\frac{\kappa + 1}{\ln(1/\overline{p_{\text{err}}})} \right] \right] N^2. \quad (\text{E7})$$

□

The next result states that any G -local Hamiltonian admits a (κ, α, τ) decomposition.

Lemma 9 (Local decompositions). Let H be an N -qubit τ -body local Hamiltonian on a graph $G = (V, E)$, and let $\tau \in O(1)$. Let $\chi(G^{(2)})$ and $\iota(G^{(2)})$ denote the chromatic and independence numbers of $G^{(2)}$ [91]. Then, H admits a (κ, α, τ) decomposition with $\kappa \leq \chi(G^{(2)}) \leq \tau^2 + 1$, $\alpha \leq \iota(G^{(2)})/N < 1$. Furthermore, any (κ, α, τ) decomposable Hamiltonian admits an on-site $(\kappa 4^{\tau}, \alpha, 1)$ decomposition.

The bounds in Lemma 9 are not necessarily tight for commuting Hamiltonians. For instance, the fully connected Ising Hamiltonian $H = -\sum_{i,j} J_{i,j} Z_i Z_j - \mu \sum_k h_k Z_k$ admits a $(1,1,1)$ decomposition, though $\chi(G^{(2)}) = N$, because it can be measured directly via single-shot on-site Z_i measurements and classical postprocessing (as it is a polynomial of the latter commuting observables). For the Hamiltonians in architectures I–III, we also find much tighter bounds.

Lemma 10 (Local decompositions in the architectures)

Let H_a be the N -qubit Hamiltonian in architecture $a \in \{\text{I, II, III}\}$. Then,

- (i) H_{I} admits a $(2, \frac{5}{9}, 1)$ decomposition of the form (E5), where every H_i has on-site terms and DO symmetry.
- (ii) H_{II} admits a $(2, \frac{5}{9}, 1)$ decomposition of the form (E5), where every H_i has on-site terms and $\text{TI}_{(1,\infty)}$ symmetry.
- (iii) H_{III} admits a two-body $(2, \frac{5}{9}, 2)$ and an on-site $(32, \frac{5}{9}, 1)$ decomposition, where every H_i has $\text{TI}_{(\sqrt{2}, \sqrt{2})}$ symmetry.

Proof of Lemma 9.—We prove the existence of the (κ, α, τ) decomposition by constructing a partition $V = \bigcup_{i=1}^{\kappa} V_i$ such that

$$H = \sum_{i=1}^{\kappa} H_i, \quad H_i := \sum_{i \in V_i} h_i, \quad (\text{E8})$$

where the terms in $\{h_i \in H_i\}$ are τ -body terms by construction and have nonoverlapping support for all $i \in [1, \kappa]$ (and, hence, can be simultaneously measured). First, $\chi(G^{(2)})$ is the minimal number of classes in any vertex coloring of $G^{(2)}$ (i.e., a vertex partition where no pair of adjacent vertices falls in the same class). Furthermore, two vertices $v_1, v_2 \in V^{(2)}$ are adjacent in $G^{(2)}$ iff they are neighbors or next neighbors in G . Letting $V = \bigcup_{i=1}^{\chi(G^{(2)})} V_i$ be a minimal vertex coloring of $G^{(2)}$, we obtain a decomposition of the form (E8) with $\kappa \leq \chi(G^{(2)})$. Last, picking $\alpha := \max_i |V_i|/N$, we get $\alpha \leq \iota(G^{(2)})/N$ since $G^{(2)}$ is connected.

The existence of the $(\kappa 4^{\deg(G)}, \alpha, 1)$ decomposition now follows by expanding each term h_i in every H_i in Eq. (E8) in a product basis

$$\mathcal{A}_i := \left\{ \bigotimes_{j \in \partial(i)} A(x_j)_j A(x_j) \in \mathcal{A} \right\}, \quad (\text{E9})$$

where $\mathcal{A} = \{A(\mu)\}_{\mu=1}^4$ is some Hermitian single-qubit operator basis (e.g., the standard Pauli matrices). By picking a fixed ordering of every set $\partial(j)$, this lets us write each H_i as a sum of 4^{τ} Hermitian operators $\{H_{i,x}, x = (x_1, \dots, x_{\tau}), x_i = [1, 4]\}_x$, $H_{i,x} := \sum_{j \in V_i} \alpha_{j,x} \bigotimes_{k \in \partial(j)} A(x_k)_k$, the energy distribution of

which can be sampled from on-site $A(x_i)$ measurements.

Finally, we upper bound $\chi(G^{(2)})$. Let $\deg(G)$ denote the maximum vertex degree of G . Since every vertex $v \in V^{(2)}$ has at most $\deg(G)$ neighbors and at most $\deg(G)(\deg(G) - 1)$ next neighbors, it follows that $\deg(G^{(2)}) \leq \deg(G)^2$. By Brook's theorem [91], $\chi(G^{(2)}) \leq \deg(G)^2 + 1$. Last, we use that $\tau = \deg(G)$ since the τ is, by definition, the maximum vertex degree of the interaction graph. □

Proof of Lemma 10.—For $H_{\text{I,II}}$, we first split the Hamiltonian terms in Eq. (E1) into two groups (“even” and “odd”), using a bicoloring of the N -site square lattice, and set $H = H_{\text{even}} + H_{\text{odd}}$. The terms $\{h_i\}_i$ of H_{even} (H_{odd}) are products of $X_{\tau_i,i}, Z_j$ on-site factors. Because we use a 2-coloring, the on-site factor list associated with two distinct h_i, h_j terms contains at most two overlapping on-site pairs of the form $(Z_k, Z_k), (Z'_k, Z'_k)$. Hence, overlapping terms are identical and can be measured jointly. This allows us to measure H_{even} (H_{odd}) via a parallel measurement of all on-site factors and classical postprocessing. This yields a (κ, α, τ) decomposition with $\kappa = 2$ and $\tau = 1$. Furthermore, the largest component V_{max} of a square lattice 2-coloring has $N/2$ vertices for even N , and $(N + 1)/2$ otherwise. Hence, we can pick $\alpha = |V_{\text{max}}|/N \leq 1/2(1 + 1/N) \leq 5/9$, where we use the fact that the smallest odd value of N is 9. The same approach leads to a $(2, 5/9, 2)$ decomposition for H_{III} , Eq. (E2), where $H_{\text{even}}, H_{\text{odd}}$ are sums of products of on-site terms supported on the primitive lattice and two-body terms acting on dangling bonds. This leads to a $(32, 5/9, 1)$ decomposition by expanding the two-body terms on a local basis, as in the proof of Lemma 9. Finally, all new Hamiltonians inherit the symmetry of their corresponding parent Hamiltonian by construction. \square

3. Certification protocol

Finally, we describe a quadratic-time weak-membership certification protocol for ground states of nondegenerate, gapped, G -local, τ -body Hamiltonians with constant τ and, in general, any (κ, α, τ) -measurable Hamiltonians (Definition 7). By virtue of Lemma 10, the protocol can be applied to efficiently certify the final-state preparation of our proposed quantum architectures I–III. We describe the protocol for the latter class since we know any G -local Hamiltonian is of that form (Lemmas 9 and 10). The protocol is simply a parallelized version of the one in Ref. [51].

Definition 11 (Weak-membership quantum-state certification [92]). Let $F_T > 0$ be a threshold fidelity and $0 < p_{\text{err}} < 1$ a maximal failure probability. A test that takes as an input a classical description of ρ_0 and copies of a preparation of ρ_p , and outputs “reject” or “accept” is a weak-membership certification test if, with high probability $p_{\text{succ}} \geq 1 - p_{\text{err}}$, it rejects every ρ_p for which $F(\rho_p, \rho_0) \leq F_T$ and accepts every ρ_p for which $F(\rho_p, \rho_0) \geq F_T + \delta$ for some fidelity gap $\delta > 0$.

Protocol 1 (Certification of (κ, α, τ) -decomposable Hamiltonians) The protocol receives a description of a nondegenerate gapped Hamiltonian H that admits a (κ, α, τ) decomposition of the form (E5), which is given to us, and performs the following steps:

- (1) Arthur chooses a threshold fidelity $F_T < 1$, maximal failure probability $1 > p_{\text{err}} > 0$, and an error $\epsilon \leq (1 - F_T)/2$.

- (2) Arthur asks Merlin to prepare a sufficient number of copies of the ground state ρ_0 of H .
- (3) Arthur performs m energy measurements for each Hamiltonian term H_i on distinct copies of the state ρ_p to determine an estimate E^* of the expectation value $\sum_i \text{tr}[\rho_p H_i]$, with m given by expression (E6). Each H_i is measured by a single-shot circuit of τ -local observables and classical postprocessing.
- (4) From the estimate E^* , he obtains an estimate F_{min}^* of lower bound $F_{\text{min}} = 1 - \langle H \rangle_{\rho_p} / \Delta$ [51] on the fidelity $F = F(\rho_p, \rho_0)$ such that $F_{\text{min}}^* \in [F_{\text{min}} - \epsilon, F_{\text{min}} + \epsilon]$ with probability of at least $1 - p_{\text{err}}$.
- (5) If $F_{\text{min}}^* < F_T + \epsilon$, he rejects; otherwise, he accepts.

Lemma 12 (Weak-membership certification). Let H be an N -qubit, nondegenerate, gapped (κ, α, τ) -decomposable Hamiltonian with known ground-state energy E_0 , gap Δ , and interaction strength $J = \max_\lambda \|h_\lambda\|$, and let E_0, Δ^{-1}, J be upper bounded by a constant. Then, Protocol 1 is a weak-membership certification test, in the sense of Ref. [51], with fidelity gap

$$\delta = (1 - F_T) \left(1 - \frac{\Delta}{\|H\|} \right) + \frac{2\epsilon\Delta}{\|H\|}, \text{ and requires}$$

$$m \geq \left[\frac{\alpha^2 \kappa^2 J^2}{2\Delta^2 \epsilon^2} \ln \left[-\frac{\kappa + 1}{\ln(1 - p_{\text{err}})} \right] \right] N^2 \quad (\text{E10})$$

repetitions to determine the expectation value $\langle H \rangle_{\rho_p} = \text{tr}[H\rho_p]$.

In combination with Lemma 10, it follows that one can efficiently certify the final-state preparation of our architectures I–III, if the latter are at least $1/N$ close to the target state in fidelity since $\|H\| \sim N$ and Δ is larger than a constant by construction of the parent Hamiltonians $H_{\text{I–III}}$.

Proof of Lemma 12.—The proof is analogous to that of Proposition 1 in Ref. [51] if we substitute Protocol 1 (Lemma 1) therein with our Protocol 1 (our Lemma 1) in this appendix. We refer the reader to Ref. [51] for details. \square

APPENDIX F: ALTERNATIVE WEAKER FORMS OF CONJECTURES 1 AND 2

In this appendix, we briefly discuss how our main result, Theorem 1, holds given even weaker versions of Conjectures 1 and 2. First, in Conjecture 1, it suffices for our purposes that the polynomial hierarchy does not collapse to its third level, instead of being infinite. Second, in Conjecture 2, we do not need the problem of approximating Ising partition functions to be #P-hard on average. Instead, it is enough that this problem is not in the complexity class BPP^{NP} , which is contained in the third level of the polynomial hierarchy (and would be in the second if the widely believed conjecture $\text{P} = \text{BPP}$ [34] holds). Note how this would be in striking contrast with our

hardness Lemma 2 since the latter says that an oracle to solve the worst-case version of the same problem would allow us to solve *all* problems in *all* levels of the polynomial hierarchy. Third, if this weaker form of Conjecture 2 holds, then Conjecture 1 is obviously not needed in the proof of Theorem 1. We have also mentioned in the main text that stating Conjecture 2 in terms of relative errors—which is the approach followed here and in Refs. [16,18,20,22]—is somewhat more natural than stating it in terms of additive ones, as in Ref. [19]. To illustrate the difference, note that there exist quantum algorithms for approximating (normalized) Ising partition functions up to polynomially small additive errors [93,94], while the latter are #P-hard to approximate up to polynomially small and even constant relative ones.

-
- [1] P. W. Shor, *Algorithms for Quantum Computation: Discrete Logarithms and Factoring*, in *Proceedings of the 35th Annual Symposium Foundations of Computer Science* (IEEE, New York, 1994), pp. 124–134.
- [2] S. Lloyd, *Universal Quantum Simulators*, *Science* **273**, 1073 (1996).
- [3] P. W. Shor, *Fault-Tolerant Quantum Computation*, in *Proceedings of the 37th Annual Symposium Foundations of Computer Science* (IEEE, New York, 1996), pp. 56–65.
- [4] E. Knill, R. Laflamme, and W. Zurek, *Threshold Accuracy for Quantum Computation*, [arXiv:quant-ph/9610011](https://arxiv.org/abs/quant-ph/9610011).
- [5] D. Aharonov and M. Ben-Or, *Fault-Tolerant Quantum Computation with Constant Error Rate*, *SIAM J. Comput.* **38**, 1207 (2008).
- [6] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, *Surface Codes: Towards Practical Large-Scale Quantum Computation*, *Phys. Rev. A* **86**, 032324 (2012).
- [7] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, Simon J. D., C. Wunderlich, and W. K. Hensinger, *Blueprint for a Microwave Trapped Ion Quantum Computer*, *Sci. Adv.* **3**, e1601540 (2017).
- [8] J. Preskill, *Quantum Computing and the Entanglement Frontier*, [arXiv:1203.5813](https://arxiv.org/abs/1203.5813).
- [9] I. Bloch, J. Dalibard, and S. Nascimbene, *Quantum Simulations with Ultracold Quantum Gases*, *Nat. Phys.* **8**, 267 (2012).
- [10] J. Eisert, M. Friesdorf, and C. Gogolin, *Quantum Many-Body Systems out of Equilibrium*, *Nat. Phys.* **11**, 124 (2015).
- [11] S. Trotzky, Y.-A. Chen, A. Flesch, I. P. McCulloch, U. Schollwöck, J. Eisert, and I. Bloch, *Probing the Relaxation Towards Equilibrium in an Isolated Strongly Correlated One-Dimensional Bose Gas*, *Nat. Phys.* **8**, 325 (2012).
- [12] J.-Y. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, *Exploring the Many-Body Localization Transition in Two Dimensions*, *Science* **352**, 1547 (2016).
- [13] S. Braun, M. Friesdorf, S. S. Hodgman, M. Schreiber, J. P. Ronzheimer, A. Riera, M. del Rey, I. Bloch, J. Eisert, and U. Schneider, *Emergence of Coherence and the Dynamics of Quantum Phase Transitions*, *Proc. Natl. Acad. Sci. U.S.A.* **112**, 3641 (2015).
- [14] P. Bordia, H. Lüschen, S. Scherg, S. Gopalakrishnan, M. Knap, U. Schneider, and I. Bloch, *Probing Slow Relaxation and Many-Body Localization in Two-Dimensional Quasiperiodic Systems*, *Phys. Rev. X* **7**, 041047 (2017).
- [15] A. Acin, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, F. Jelezko, S. Kuhr, M. Lewenstein, M. F. Riedel, P. O. Schmidt, R. Thew, A. Wallraff, I. Walmsley, and F. K. Wilhelm, *The European Quantum Technologies Roadmap*, [arXiv:1712.03773](https://arxiv.org/abs/1712.03773).
- [16] S. Aaronson and A. Arkhipov, *The Computational Complexity of Linear Optics*, *Theory Comput.* **9**, 143 (2013).
- [17] 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).
- [18] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, J. M. Martinis, and H. Neven, *Characterizing Quantum Supremacy in Near-Term Devices*, [arXiv:1608.00263](https://arxiv.org/abs/1608.00263).
- [19] X. Gao, S.-T. Wang, and L.-M. Duan, *Quantum Supremacy for Simulating a Translation-Invariant Ising Spin Model*, *Phys. Rev. Lett.* **118**, 040502 (2017).
- [20] M. J. Bremner, A. Montanaro, and D. J. Shepherd, *Achieving Quantum Supremacy with Sparse and Noisy Commuting Quantum Computations*, *Quantum* **1**, 8 (2017).
- [21] B. M. Terhal and D. P. DiVincenzo, *Adaptive Quantum Computation, Constant Depth Quantum Circuits and Arthur-Merlin Games*, *Quantum Inf. Comput.* **4**, 134 (2004).
- [22] M. J. Bremner, R. Jozsa, and D. J. Shepherd, *Classical Simulation of Commuting Quantum Computations Implies Collapse of the Polynomial Hierarchy*, *Proc. R. Soc. A* **467**, 459 (2011).
- [23] J. I. Cirac and P. Zoller, *Goals and Opportunities in Quantum Simulation*, *Nat. Phys.* **8**, 264 (2012).
- [24] M. Endres, M. Cheneau, T. Fukuhara, C. Weitenberg, P. Schau, C. Gross, L. Mazza, M. C. Bauls, L. Pollet, I. Bloch *et al.*, *Single-Site- and Single-Atom-Resolved Measurement of Correlation Functions*, *Appl. Phys. B* **113**, 27 (2013).
- [25] O. Mandel, M. Greiner, A. Widera, T. Rom, T. W. Hänsch, and I. Bloch, *Controlled Collisions for Multiparticle Entanglement of Optically Trapped Atoms*, *Nature (London)* **425**, 937 (2003).
- [26] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, *A Quantum Gas Microscope for Detecting Single Atoms in a Hubbard-Regime Optical Lattice*, *Nature (London)* **462**, 74 (2009).
- [27] D. Shepherd and M. J. Bremner, *Temporally Unstructured Quantum Computation*, *Proc. R. Soc. A* **465**, 1413 (2009).
- [28] M. J. Hoban, J. J. Wallman, H. Anwar, N. Usher, R. Raussendorf, and D. E. Browne, *Measurement-Based Classical Computation*, *Phys. Rev. Lett.* **112**, 140505 (2014).
- [29] R. Raussendorf and H. J. Briegel, *A One-Way Quantum Computer*, *Phys. Rev. Lett.* **86**, 5188 (2001).
- [30] S. Bravyi and D. Gosset, *Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates*, *Phys. Rev. Lett.* **116**, 250501 (2016).
- [31] This also implies that the entire resource does not necessarily have to be prepared in a single preparation but can be sequentially prepared.

- [32] R. Jozsa and M. Van Den Nest, *Classical Simulation Complexity of Extended Clifford Circuits*, *Quantum Inf. Comput.* **14**, 633 (2014).
- [33] B. Fefferman and C. Umans, *On the Power of Quantum Fourier Sampling*, in *Proceedings of the 11th Conference on the Theory of Quantum Computation, Communications, and Cryptography*, Vol. 61 (Schloss Dagstuhl Publishing, Germany, 2016), pp. 1:1–1:19.
- [34] S. Aaronson, $P \neq NP?$ in *Open Problems in Mathematics* (Springer, New York, 2016).
- [35] L. Fortnow, *Beyond NP: The Work and Legacy of Larry Stockmeyer*, in *Proceedings of the 37th Annual Symposium on Theory of Computing* (ACM, Baltimore, MD, 2005), pp. 120–127.
- [36] R. M. Karp and R. J. Lipton, *Some Connections Between Nonuniform and Uniform Complexity Classes*, in *Proceedings of the 12th Annual Symposium on Theory of Computing* (ACM, Baltimore, MD, 1980), pp. 302–309.
- [37] L. G. Valiant, *The Complexity of Computing the Permanent*, *Theor. Comput. Sci.* **8**, 189 (1979).
- [38] R. Raussendorf, D. E. Browne, and H. J. Briegel, *Measurement-Based Quantum Computation on Cluster States*, *Phys. Rev. A* **68**, 022312 (2003).
- [39] R. Raussendorf, *Quantum Computation via Translation-Invariant Operations on a Chain of Qubits*, *Phys. Rev. A* **72**, 052301 (2005).
- [40] A. Mantri, T. F. Demarie, and J. F. Fitzsimons, *Universality of Quantum Computation with Cluster States and (X,Y)-plane Measurements*, *Sci. Rep.* **7**, 42861 (2017).
- [41] C. E. Porter and R. G. Thomas, *Fluctuations of Nuclear Reaction Widths*, *Phys. Rev.* **104**, 483 (1956).
- [42] F. Haake, *Quantum Signatures of Chaos*, Physics and Astronomy Online Library (Springer, New York, 2001).
- [43] J. Emerson, Y. S. Weinstein, M. Saraceno, S. Lloyd, and D. G. Cory, *Pseudo-random Unitary Operators for Quantum Information Processing*, *Science* **302**, 2098 (2003).
- [44] J. Emerson, E. Livine, and S. Lloyd, *Convergence Conditions for Random Quantum Circuits*, *Phys. Rev. A* **72**, 060302 (2005).
- [45] W. G. Brown, Y. S. Weinstein, and L. Viola, *Quantum Pseudorandomness from Cluster-State Quantum Computation*, *Phys. Rev. A* **77**, 040303 (2008).
- [46] S. Aaronson and L. Chen, *Complexity-Theoretic Foundations of Quantum Supremacy Experiments*, [arXiv:1612.05903](https://arxiv.org/abs/1612.05903).
- [47] A. P. Lund, M. J. Bremner, and T. C. Ralph, *Quantum Sampling Problems, Boson Sampling and Quantum Supremacy*, *Quantum Inf. Comput.* **3**, 15 (2017).
- [48] C. Brand, H. Dell, and M. Roth, *Fine-Grained Dichotomies for the Tutte Plane and Boolean #CSP*, in *11th International Symposium on Parameterized and Exact Computation*, Leibniz International Proceedings in Informatics, Vol. 63 (Schloss Dagstuhl Publishing, Germany, 2017), pp. 9:1–9:14.
- [49] R. Curticapean, *Block Interpolation: A Framework for Tight Exponential-Time Counting Complexity*, in *Proceedings of the 42nd International Colloquium of Automata, Languages, and Programming* (Springer, New York, 2015), pp. 380–392.
- [50] M. Schwarz and M. Van den Nest, *Simulating Quantum Circuits with Sparse Output Distributions*, [arXiv:1310.6749](https://arxiv.org/abs/1310.6749).
- [51] D. Hangleiter, M. Kliesch, M. Schwarz, and J. Eisert, *Direct Certification of a Class of Quantum Simulations*, *Quant. Sci. Tech.* **2**, 015004 (2017).
- [52] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, *Efficient Quantum State Tomography*, *Nat. Commun.* **1**, 149 (2010).
- [53] J. Miller, S. Sanders, and A. Miyake, *Quantum Supremacy in Constant-Time Measurement-Based Computation: A Unified Architecture for Sampling and Verification*, *Phys. Rev. A* **96**, 062320 (2017).
- [54] F. Verstraete and J. I. Cirac, *Valence-Bond States for Quantum Computation*, *Phys. Rev. A* **70**, 060302 (2004).
- [55] D. Gross and J. Eisert, *Novel Schemes for Measurement-Based Quantum Computation*, *Phys. Rev. Lett.* **98**, 220503 (2007).
- [56] D. Poulin, *Stabilizer Formalism for Operator Quantum Error Correction*, *Phys. Rev. Lett.* **95**, 230504 (2005).
- [57] M. Hein, J. Eisert, and H. J. Briegel, *Multi-particle Entanglement in Graph States*, *Phys. Rev. A* **69**, 062311 (2004).
- [58] D. Jaksch, H.-J. Briegel, J. I. Cirac, C. W. Gardiner, and P. Zoller, *Entanglement of Atoms via Cold Controlled Collisions*, *Phys. Rev. Lett.* **82**, 1975 (1999).
- [59] A. Widera, F. Gerbier, S. Fölling, T. Gericke, O. Mandel, and I. Bloch, *Coherent Collisional Spin Dynamics in Optical Lattices*, *Phys. Rev. Lett.* **95**, 190405 (2005).
- [60] S. Nascimbène, Y.-A. Chen, M. Atala, M. Aidelsburger, S. Trotzky, B. Paredes, and I. Bloch, *Experimental Realization of Plaquette Resonating Valence-Bond States with Ultracold Atoms in Optical Superlattices*, *Phys. Rev. Lett.* **108**, 205301 (2012).
- [61] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner, *Probing the Superfluid-to-Mott Insulator Transition at the Single-Atom Level*, *Science* **329**, 547 (2010).
- [62] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schaub, T. Fukuhara, I. Bloch, and S. Kuhr, *Single-spin Addressing in an Atomic Mott Insulator*, *Nature (London)* **471**, 319 (2011).
- [63] C. Gross, private communication.
- [64] M. Schreiber, S. S. Hodgman, P. Bordia, H. P. Lüschen, M. H. Fischer, R. Vosk, E. Altman, U. Schneider, and I. Bloch, *Observation of Many-Body Localization of Interacting Fermions in a Quasirandom Optical Lattice*, *Science* **349**, 842 (2015).
- [65] A. Singha, M. Gibertini, B. Karmakar, S. Yuan, M. Polini, G. Vignale, M. I. Katsnelson, A. Pinczuk, L. N. Pfeiffer, K. W. West, and V. Pellegrini, *Two-Dimensional Mott-Hubbard Electrons in an Artificial Honeycomb Lattice*, *Science* **332**, 1176 (2011).
- [66] T. Jacqmin, I. Carusotto, I. Sagnes, M. Abbarchi, D. D. Solnyshkov, G. Malpuech, E. Galopin, A. Lemaître, J. Bloch, and A. Amo, *Direct Observation of Dirac Cones and a Flatband in a Honeycomb Lattice for Polaritons*, *Phys. Rev. Lett.* **112**, 116402 (2014).
- [67] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis,

- Superconducting Quantum Circuits at the Surface Code Threshold for Fault Tolerance*, *Nature (London)* **508**, 500 (2014).
- [68] A. A. Houck, H. E. Türeci, and J. Koch, *On-Chip Quantum Simulation with Superconducting Circuits*, *Nat. Phys.* **8**, 292 (2012).
- [69] M. R. Geller, J. M. Martinis, A. T. Sornborger, P. C. Stancil, E. J. Pritchett, H. You, and A. Galiatdinov, *Universal Quantum Simulation with Prethreshold Superconducting Qubits: Single-Excitation Subspace Method*, *Phys. Rev. A* **91**, 062309 (2015).
- [70] On the one hand, configurable lattice systems of two-dimensional arrays of individually trapped ions can be conceived [71]. Probably more feasible, on the other hand, are flattened architectures, in which the 2D setting is traded for suitable long-ranged quantum gates.
- [71] M. Mielenz, H. Kalis, M. Wittemer, F. Hakeberg, R. Schmied, M. Blain, P. Maunz, D. Leibfried, U. Warring, and T. Schaetz, *Freely Configurable Quantum Simulator Based on a Two-Dimensional Array of Individually Trapped Ions*, arXiv:1512.03559.
- [72] S. Aaronson, *Quantum Computing, Postselection, and Probabilistic Polynomial-Time*, *Proc. R. Soc. A* **461**, 3473 (2005).
- [73] L. Stockmeyer, *On Approximation Algorithms for #P*, *SIAM J. Comput.* **14**, 849 (1985).
- [74] S. Toda, *PP Is as Hard as the Polynomial-Time Hierarchy*, *SIAM J. Comput.* **20**, 865 (1991).
- [75] M. A. Nielsen, *Conditions for a Class of Entanglement Transformations*, *Phys. Rev. Lett.* **83**, 436 (1999).
- [76] P. O. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan, *A New Universal and Fault-Tolerant Quantum Basis*, *Inf. Proc. Lett.* **75**, 101 (2000).
- [77] R. Raussendorf, *Quantum Cellular Automaton for Universal Quantum Computation*, *Phys. Rev. A* **72**, 022301 (2005).
- [78] A. M. Childs, D. W. Leung, and M. A. Nielsen, *Unified Derivations of Measurement-Based Schemes for Quantum Computation*, *Phys. Rev. A* **71**, 032318 (2005).
- [79] X. Zhou, D. W. Leung, and I. L. Chuang, *Methodology for Quantum Logic Gate Construction*, *Phys. Rev. A* **62**, 052316 (2000).
- [80] K. Fujii and T. Morimae, *Quantum Commuting Circuits and Complexity of Ising Partition Functions*, *New J. Phys.* **19**, 033003 (2017).
- [81] L. A. Goldberg and H. Guo, *The Complexity of Approximating Complex-Valued Ising and Tutte Partition Functions*, arXiv:1409.5627.
- [82] M. Soeken, D. M. Miller, and R. Drechsler, *Quantum Circuits Employing Roots of the Pauli Matrices*, *Phys. Rev. A* **88**, 042322 (2013).
- [83] A. W. Harrow and R. A. Low, *Random Quantum Circuits Are Approximate 2-Designs*, *Commun. Math. Phys.* **291**, 257 (2009).
- [84] H. Kim and D. A. Huse, *Ballistic Spreading of Entanglement in a Diffusive Nonintegrable System*, *Phys. Rev. Lett.* **111**, 127205 (2013).
- [85] P. Hosur, X.-L. Qi, D. A. Roberts, and B. Yoshida, *Chaos in Quantum Channels*, *J. High Energy Phys.* **02** (2016) 4.
- [86] 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).
- [87] D. Wecker and K. M. Svore, *LIQ|>: A Software Design Architecture and Domain-Specific Language for Quantum Computing*, arXiv:1402.4467.
- [88] M. Hein, W. Dür, J. Eisert, R. Raussendorf, M. Van den Nest, and H.-J. Briegel, *Entanglement in Graph States and Its Applications*, in *Quantum Computers, Algorithms and Chaos*, International School of Physics, Enrico Fermi (IOS Press, Amsterdam, 2006), arXiv:quant-ph/0602096.
- [89] S. Bravyi and R. Raussendorf, *Measurement-based Quantum Computation with the Toric Code States*, *Phys. Rev. A* **76**, 022304 (2007).
- [90] M. Van den Nest, W. Dür, and H. J. Briegel, *Completeness of the Classical 2D Ising Model and Universal Quantum Computation*, *Phys. Rev. Lett.* **100**, 110501 (2008).
- [91] J. A. Bondy, *Graph Theory with Applications* (Elsevier Science, Oxford, England, 1976).
- [92] L. Aolita, C. Gogolin, M. Kliesch, and J. Eisert, *Reliable Quantum Certification of Photonic State Preparations*, *Nat. Commun.* **6**, 8498 (2015).
- [93] G. De las Cuevas, W. Duer, M. Van den Nest, and M. A. Martin-Delgado, *Quantum Algorithms for Classical Lattice Models*, *New J. Phys.* **13**, 093021 (2011).
- [94] A. Matsuo, K. Fujii, and N. Imoto, *Quantum Algorithm for an Additive Approximation of Ising Partition Functions*, *Phys. Rev. A* **90**, 022304 (2014).