Efficiently verifiable quantum advantage on near-term analog quantum simulators

Zhenning Liu[1,2,3], Dhruv Devulapalli[1,2], Dominik Hangleiter[1], Yi-Kai Liu[1,4],
Alicia J. Kollár[2,5], Alexey V. Gorshkov[1,2], and Andrew M. Childs[1,3,6]

[1] Joint Center for Quantum Information and Computer Science, NIST/University of Maryland, College Park, Maryland 20742, USA
[2] Joint Quantum Institute, NIST/University of Maryland, College Park, Maryland 20742, USA
[3] Department of Computer Science, University of Maryland, College Park, Maryland 20742, USA
[4] Applied and Computational Mathematics Division, National Institute of Standards and Technology (NIST), Gaithersburg, Maryland 20899, USA
[5] Department of Physics, University of Maryland, College Park, Maryland 20742, USA
[6] Institute for Advanced Computer Studies, University of Maryland, College Park, Maryland 20742, USA

Existing schemes for demonstrating quantum computational advantage are subject to various practical restrictions, including the hardness of verification and challenges in experimental implementation. Meanwhile, analog quantum simulators have been realized in many experiments to study novel physics. In this work, we propose a quantum advantage protocol based on *single-step Feynman-Kitaev* verification of an analog quantum simulation, in which the verifier need only run an $O(\lambda^2)$-time classical computation, and the prover need only prepare $O(1)$ samples of a history state and perform $O(\lambda^2)$ single-qubit measurements, for a security parameter λ. We also propose a near-term feasible strategy for honest provers and discuss potential experimental realizations.

I. INTRODUCTION

A. Background & Motivation

Quantum computers offer the promise of executing some computational tasks exponentially faster than classical computers. This suggests a violation of the extended Church-Turing thesis, which says that any physically realizable model of computation can be efficiently simulated by a classical Turing machine. Indeed, quantum computers were originally proposed as a means of simulating quantum mechanical systems [1], a task considered classically hard. There has been much progress toward identifying classically difficult problems that quantum computers can solve efficiently, such as integer factorization [2], simulating Hamiltonian dynamics [3–5], and extracting information about solutions of high-dimensional linear systems [6].

A significant milestone for the field of quantum computing is the first demonstration that a quantum device can perform computational tasks that a classical device with comparable resources cannot. This milestone has been called quantum supremacy [7, 8], quantum advantage, or a proof of quantumness [9], and has instigated numerous theoretical proposals and experimental efforts. However, there remain formidable technological challenges to building quantum computers, requiring both theoretical and experimental progress in architecture design, fault tolerance, and control. Various proposals for quantum advantage have addressed these challenges in different ways, by making trade-offs between ease of experimental demonstration, ease of verification, security guarantees, and practical applications.

Analog quantum simulation [10], i.e., using one many-body quantum system to simulate another, is a natural approach to demonstrating quantum advantage. By building quantum systems with tunable (but perhaps non-universal) Hamiltonians, one can emulate a large class of Hamiltonians that may be difficult to simulate classically. Since it directly encodes hard problems into controllable quantum systems, analog simulation arguably mitigates many of the issues faced by digital approaches [11, 12]. Furthermore, analog simulation avoids Trotter error and other sources of algorithmic error in digital quantum simulation [13, 14]. Indeed, analog simulations of systems with hundreds of qubits have already been performed [15].

A major challenge for both quantum simulation and more general forms of quantum computation is the difficulty of verifying the correctness of a quantum process. There have been several proposals to verify digital quantum computation [16, 17] based on the *Feynman-Kitaev circuit-to-Hamiltonian mapping* [18], but such protocols are neither designed for analog quantum simulation nor practical on near-term analog quantum devices. Previous work towards verifying analog simulation has suggested approaches such as cross-platform verification [19, 20], Hamiltonian learning [20], and performing a Loschmidt echo [20–22]. Unlike protocols for digital verification, these approaches can be spoofed by dishonest or inaccurate quantum simulators, and therefore cannot be used to demonstrate quantum advantage in a sound, efficiently verifiable way. A step toward verified analog simulation is made in [23], in which the verifier measures the energy of a *parent Hamiltonian* of the output state of analog quantum simulation. However, all these works require a significant number of samples of the simulator's state to certify it.

B. Our Contribution

In this paper, by combining a *single-step* Feynman-Kitaev encoding and the scheme of Bermejo-Vega et al. [24], we propose a novel quantum advantage protocol with reduced resource requirements, where a verifier ca-

pable of polynomial-time classical computation can verify the result by asking the prover to perform *trusted measurements* on a *constant* number of copies of a state. We also present a strategy for the honest prover and argue that it is feasible on near-term devices.

The Protocol. Our protocol involves interaction between a polynomial-time classical *verifier*, and a quantum *prover* who can do polynomial-time quantum computation, although we present a strategy for an honest prover who must only perform analog quantum simulation and some limited additional operations. In our protocol, the prover is capable of single-qubit *trusted measurements*, which means that he performs the correct single-qubit measurements as instructed by the verifier with error rate $\epsilon = O(1/n)$ (with n the number of qubits), and reports the outcome honestly. We also allow a polynomial amount of classical communication in both directions.

Our protocol still works without the assumption of trusted measurements if the prover can send polynomial-size quantum states to the verifier, and the verifier can perform single-qubit measurements, as in the notion of a non-interactive QPIP_1 protocol defined by Aharonov et al. [25] (where QPIP stands for *quantum prover interactive proof*).

Definition 1 (QPIP_k protocol (simplified))**.** *An interactive proof for a language $\mathcal{L}$ is said to be QPIP_k if the prover is a* BQP *machine, the verifier is a hybrid* BQP-BPP *machine that can process at most k qubits at a time, and quantum states of k qubits can be transmitted from the prover to the verifier.*

However, as reliably sending quantum states is unlikely to be feasible in the near term, we focus on the former model.

Prover's Model of Computation. We also give an experimentally practical strategy for honest provers. The strategy is specifically designed for near-term machines that are not capable of fully digital quantum computation, but are slightly more powerful than *analog* quantum simulation, a popular notion that is often not clearly defined. In our work, we define a *mostly analog* model of computation, its *commuting* version, and its extension with a global CZ gate, which we argue are feasible for near-term experiments.

Definition 2 (Mostly analog quantum computation)**.** *A model of quantum computation involving n qubits is called* mostly analog *if all the following conditions hold. (1) The system can evolve under a* time-independent *2-body Hamiltonian H containing* poly(n) *Pauli terms for time $T = \mathrm{poly}(n)$. (2) $O(1)$ alternations between the evolution under H and single-qubit rotations can be performed. (3) Measurements can only be performed once at the end of the whole process.*

Note that condition (2) distinguishes this model from common notions of analog quantum computation, as it may require a degree of control not always available to analog quantum simulators. Despite being mostly analog, the above model of computation is even capable of BQP-complete quantum computations [26]. We introduce a weaker model where the 2-local Hamiltonian H must also be *commuting*, which means that all Pauli terms must commute with each other.

Definition 3 (Mostly analog commuting quantum computation)**.** *A mostly analog model of computation is called* commuting *if H is a commuting Hamiltonian.*

Even a mostly analog commuting quantum device can solve some classically intractable problems [24]. We focus on an even more restricted model that should be easier to realize, in which the Hamiltonian H is a specific commuting Hamiltonian containing only nearest-neighbor Z operators, as discussed further below.

We also assume the ability to perform a globally controlled CZ gate. This arguably makes our model less analog, but it plays a key role in developing a sample-efficient protocol to verify the solutions given by the device, and it can potentially be realized using experimental capabilities that have already been demonstrated [27, 28], as we discuss in Section III B.

Definition 4 (Mostly analog + GCZ commuting quantum computation)**.** *A mostly analog commuting model of computation is called* mostly analog + GCZ *if the system also contains a quantum degree of freedom (e.g., a qubit) that can serve as a global control for all of the qubits, such that one can apply—only $O(1)$ times—a global CZ gate that is controlled by the degree of freedom and acts on all of the qubits. Here* GCZ *stands for global CZ.*

The Classically Hard Problem. In the protocol, the verifier asks the prover to solve a classically hard problem based on Hamiltonian evolution. The prover generates a quantum state but is not trusted to do so correctly. However, the prover is trusted to honestly measure this state to generate a classical witness. The verifier checks this witness to determine if the problem has been successfully solved. If so, then quantum advantage has been demonstrated.

Instead of considering a general quantum circuit, we aim to demonstrate quantum advantage by verifying a specific analog quantum simulation performed on a mostly analog + GCZ commuting machine. The simulation is motivated by the class of IQP (instantaneous quantum polynomial-time) circuits [29, 30], in which all quantum gates are commuting (and thus interchangeable in time). Despite this strong restriction, IQP circuits are believed to be hard to simulate classically [30, 31]. Furthermore, Bermejo-Vega et al. [24] presented a concrete scheme to show quantum speedup on an analog simulator by running a specific unit-time Hamiltonian evolution. The Hamiltonian includes only nearest-neighbor ZZ interactions and local Z terms (a form that we call a *$(ZZ + Z)$-type Hamiltonian*) on a 2-dimensional square

lattice:

$$\sum_{\{i,j\}\in \mathrm{NN}} \frac{\pi}{4} Z_i Z_j - \sum_{i=1}^{n} \frac{\pi}{4} Z_i, \tag{1}$$

where NN denotes the set of edges connecting nearest-neighbor qubits. The qubits are randomly initialized in either $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ or $\frac{1}{\sqrt{2}}(|0\rangle + e^{-\mathrm{i}\pi/4}|1\rangle)$. Bermejo-Vega et al. [24] and Ringbauer et al. [32] prove that a classical computer cannot efficiently sample from the output distribution of X-basis measurements on the above system within total variation distance (TVD) 0.292, under plausible computational assumptions that we review in Appendix S1. Moreover, since single-qubit Z_i operators commute with all Z_iZ_j operators, one can absorb the single-qubit evolution $\exp\left(\mathrm{i}\frac{\pi}{4}\sum_i Z_i\right)$ into the initial state of each qubit, so that the qubits are initialized in either $\frac{1}{2}\left[(1+\mathrm{i})|0\rangle + (1-\mathrm{i})|1\rangle\right]$ or $\frac{1}{2}\left[(1+\mathrm{i})|0\rangle + e^{-\mathrm{i}\pi/4}(1-\mathrm{i})|1\rangle\right]$, which can be prepared by single-qubit operations. Then the Hamiltonian H to be simulated contains only ZZ interaction terms:

$$H = \sum_{\{i,j\}\in \mathrm{NN}} \frac{\pi}{4} Z_i Z_j. \tag{2}$$

Main Result. We now have all the building blocks to formalize the main result. In the state-transmission scenario, we have the following theorem.

Theorem 1 (Main result—state-transmission version)**.** *There exists a classically intractable sampling problem that can be verified by a single-round* QPIP_1 *protocol where the prover runs a specific mostly analog +* GCZ *commuting quantum task* $O(1)$ *times.*

In the trusted-measurement scenario, our result is as follows.

Theorem 2 (Main result—trusted-measurement version)**.** *There exists a classically intractable sampling problem that can be verified by a single-round protocol where the classical verifier trusts the prover to perform single-qubit measurements, and the prover runs a specific mostly analog +* GCZ *commuting quantum task* $O(1)$ *times.*

Our protocol has constant *sample complexity*, i.e., it only requires the prover to generate $O(1)$ samples of an n-qubit state. This is significantly less expensive than Bermejo-Vega et al. [24], which uses $O(n^2)$ samples.

In both this work and Ref. [24], the prover is expected to perform trusted measurements (or the prover sends qubits to the verifier for her to measure), unlike proofs of quantumness (PoQs) based on trapdoor claw-free functions (TCFs) [9, 33] and quantum supremacy experiments [7, 8] based on sampling problems, which makes it difficult to compare the resource requirements. However, in all of these schemes, single-qubit measurements must be performed many times, either by the prover or the verifier. Hence the number of qubits measured is a comparable quantity.

Equivalently, without transforming the protocols, we can still compare the number of measurements in terms of the *security parameter*, whether the measurements are trusted or not. The security parameter λ is defined such that a dishonest prover without quantum computational power needs time $2^{\Omega(\lambda)}$ in order to make the verifier accept. For our protocol, the number of qubits n is quadratic in λ, just as in Bermejo-Vega et al. [24]. Under optimistic assumptions, cryptographic PoQs can probably have $n = O(\lambda)$ [33], but for most common TCFs, n scales at least quadratically with λ [9]. Since it has constant sample complexity, our protocol uses $O(\lambda^2)$ single-qubit measurements. This is better than Bermejo-Vega et al. [24], which uses $O(\lambda^3)$ measurements. Furthermore, our protocol can be verified by $O(\lambda^2)$-time classical computation, significantly below the verification cost of $O(\lambda^6)$ for Bermejo-Vega et al. [24] and presumably $\exp(\lambda)$ for quantum supremacy experiments based on sampling problems [7, 35, 36].

We summarize the comparison between our work and other quantum advantage protocols in Table I.

On the prover side, TCF-based PoQs generally require poly(λ)-depth low-noise digital quantum computation, while our honest strategy is designed for analog quantum simulators with only limited digital capabilities. This may be harder than fully analog simulation [7, 24, 37, 38], but should still be feasible in the relatively near term. Moreover, our protocol can detect—and is robust against—a specific type of phase error that happens frequently in practice. Thus we believe our work achieves a significant improvement in terms of verification efficiency for verified quantum advantage protocols, and is an easier-to-implement scheme. We provide exact threshold fidelities (independent of the system size) for the device to demonstrate quantum advantage using our scheme. We also show that when the noise is incoherent, the fidelity requirements can be further relaxed.

The remainder of this paper is organized as follows. In Section II, we describe the sample-efficient quantum advantage protocol and analyze its resource requirements. In Section III, we give the near-term strategy for honest provers and discuss potential experimental realizations. Finally, we summarize the results and discuss their implications and potential future extensions in Section IV.

II. THE QUANTUM ADVANTAGE PROTOCOL

A. The Single-Step Feynman-Kitaev Construction

Our protocol is inspired by the Feynman-Kitaev mapping [18], which converts the task of executing a quantum circuit to that of finding the ground state of an associated Hamiltonian. The Feynman-Kitaev Hamiltonian is the foundation of several verification schemes in the circuit model: if a quantum server can provide the ground

Scheme	# of Measurements	Classical Verification	Requirements for Honest Provers	Requirements for Verifiers
Cryptographic PoQs [9, 33]	poly(λ)	poly(λ)	Digital	Purely Classical
Random Circuit Sampling [7, 34]	$O(\lambda)$	exp(λ)	Digital	Purely Classical
Parent Hamiltonians [24]	$O(\lambda^6)$	$O(\lambda^6)$	Analog	Single-Qubit Measurements
This Work (State Transmission)	$O(\lambda^2)$	$O(\lambda^2)$	Mostly Analog + Global CZ	Single-Qubit Measurements
This Work (Trusted Measurements)	$O(\lambda^2)$	$O(\lambda^2)$	Mostly Analog + Global CZ + Trusted Measurements	Purely Classical

TABLE I. Comparison of demonstrations of quantum advantage. As discussed in the main text, λ denotes the security parameter.

state (the witness) to the client, then the client can verify the quantum computation by measuring its energy. Examples include the Fitzsimons et al. [16] protocol where the prover needs to perform single-qubit trusted measurements, and the Mahadev [17] protocol that works even for untrusted measurements.

Inspired by the above protocols for circuit-model computations, we consider using a simplified Feynman-Kitaev mapping to verify analog quantum simulation of the system in [24], i.e., the Hamiltonian H in Eq. (2).

We define the (single-step) *history state*

$$|\psi_{\text{hist}}\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle |\phi_{\text{in}}\rangle + |1\rangle U |\phi_{\text{in}}\rangle \right), \tag{3}$$

where $|\phi_{\text{in}}\rangle$ is the input state and U is the *propagation unitary*. The state $|\psi_{\text{hist}}\rangle$ is the ground state of the single-step Feynman-Kitaev Hamiltonian. Since we are considering quantum simulation of the ZZ-type Hamiltonian H defined in Eq. (2), we have $U = \exp(-\mathrm{i}HT)$ with $T = 1$, and $|\phi_{\text{in}}\rangle$ is the same random input state defined in the system of Bermejo-Vega et al. [24] with single-qubit Z evolution absorbed. The computationally hard sampling problem can be solved by measuring $U |\phi_{\text{in}}\rangle$ in the X basis. We use P_{ideal} to denote the ideal distribution of measurement outcomes.

The Feynman-Kitaev Hamiltonian includes a term

$$H^{\text{prop}} = \frac{1}{2} \left(I \otimes I - |1\rangle\langle 0| \otimes U - |0\rangle\langle 1| \otimes U^\dagger \right), \tag{4}$$

which ensures that the ground state encodes the correct propagation unitary U. One can easily check that H^{prop} is positive semidefinite and $H^{\text{prop}} |\psi_{\text{hist}}\rangle = 0$, so $|\psi_{\text{hist}}\rangle$ is a ground state of H^{prop}.

The other term of the Feynman-Kitaev Hamiltonian is

$$H^{\text{in}} = |0\rangle\langle 0| \otimes \left(I - \sum_i |\phi_{\text{in},i}\rangle\langle \phi_{\text{in},i}| \right), \tag{5}$$

where $|\phi_{\text{in},i}\rangle$ is the state of the ith qubit of $|\phi_{\text{in}}\rangle$. H^{in} ensures that the input state is $|\phi_{\text{in}}\rangle$. It is also positive semidefinite and satisfies $H^{\text{in}} |\psi_{\text{hist}}\rangle = 0$.

A toy version of our protocol for demonstrating quantum advantage, without any technical detail, is as follows. The verifier sends classical descriptions of H and $|\phi_{\text{in}}\rangle$ to the prover, and asks the prover to prepare N_{M} copies of the history state $\frac{1}{\sqrt{2}}(|0\rangle |\phi_{\text{in}}\rangle + |1\rangle U |\phi_{\text{in}}\rangle)$. For each copy, the verifier chooses whether to generate a sample or to verify the state, with equal probability. If she chooses to sample, then she asks the prover to measure the first qubit (i.e., the clock qubit) in the Z basis and all other qubits in the X basis, and a sample is generated if the first measurement outcome is -1 (i.e., the clock qubit is in $|1\rangle$). If the verifier chooses to verify, then she measures the energy of $H^{\text{prop}} + H^{\text{in}}$ by quantum phase estimation. Finally, if every run of quantum phase estimation returns 0, which means that the fidelity between the measured state and the perfect history state is very high (the infidelity is inverse exponential in N_{M} if $N_{\text{M}}/2$ copies are chosen for verification) and therefore the measurement outcomes are close to the desired distribution P_{ideal}, she accepts and announces all of the samples obtained. Otherwise, she rejects.

One disadvantage of the verification part of this scheme is that it can only accept devices that provide history states with exponentially small infidelity. While near-term devices will be imperfect, they might still be able to sample from classically intractable distributions. Also, experimentalists may prefer to know how well their devices are performing and whether they are making progress, but a "yes or no" result cannot provide this kind of information. Finally, measurements of $H^{\text{prop}} + H^{\text{in}}$ might be difficult, potentially requiring many measurements to determine the energy with sufficiently high precision, and quantum phase estimation is not feasible in the near term.

Therefore, inspired by the original single-step Feynman-Kitaev Hamiltonian, we propose a new verification scheme to replace the toy protocol. In the new scheme, different *parameters* are measured to lower bound the total variation distance between the sampled distribution P_{exp} and the desired distribution P_{ideal}, demonstrating quantum advantage according to [24, 32]. We also give an efficient near-term strategy for estimating those parameters.

B. Our Measurement Scheme

To begin, consider an arbitrary $(n+1)$-qubit state

$$\rho = \sum_i p_i \left|\psi_i\right\rangle \left\langle\psi_i\right|, \tag{6}$$

where $\{|\psi_i\rangle\}$ is the (unknown) eigenbasis of ρ, and p_i is the probability corresponding to $|\psi_i\rangle$. We can write

$$|\psi_i\rangle = \alpha_i |0\rangle |\phi_i\rangle + \beta_i |1\rangle |\phi_i'\rangle \tag{7}$$

where $|\phi_i\rangle$ and $|\phi_i'\rangle$ are n-qubit states and $\alpha_i, \beta_i \in \mathbb{C}$ satisfy $|\alpha_i|^2 + |\beta_i|^2 = 1$. Thus we can interpret ρ as a classical mixture of states $|\psi_i\rangle$ as above with input states $|\phi_i\rangle$ and output states $|\phi_i'\rangle$.

The first parameter to be estimated in our scheme, the *input fidelity*, is defined as

$$F_{\text{in}}(\rho) := \frac{\sum_i p_i |\alpha_i|^2 \left|\langle\phi_i|\phi_{\text{in}}\rangle\right|^2}{\sum_i p_i |\alpha_i|^2}. \tag{8}$$

This quantifies the quality of initial state preparation. It plays a similar role to $\langle H^{\text{in}}\rangle$ in the single-step Feynman-Kitaev Hamiltonian.

Another parameter is the probability of obtaining a -1 outcome when measuring the clock qubit. We call this the *probability of sampling*:

$$p_{\text{samp}} := \sum_i p_i |\beta_i|^2. \tag{9}$$

The last parameter is the *de facto* mean value of the non-Hermitian operator

$$O_{10} := |1\rangle\langle 0| \otimes U, \tag{10}$$

whose expectation value in the state ρ is

$$\text{Tr}[\rho O_{10}] = \sum_i p_i \alpha_i \beta_i^* \langle\phi_i'|U|\phi_i\rangle. \tag{11}$$

We mainly consider its squared norm, $|\text{Tr}[\rho O_{10}]|^2$. This quantity is related to the quality of propagation from $|\phi\rangle$ to $U|\phi\rangle$, so it plays a similar role to $\langle H^{\text{prop}}\rangle$ in the single-step Feynman-Kitaev Hamiltonian.

As we show in Lemmas 1 and 2, $F_{\text{in}}(\rho)$, p_{samp}, and $\text{Tr}[\rho O_{10}]$ can all be estimated by single-qubit measurements, and the precision depends only on the number of samples measured, independent of the system size. Note that O_{10} is not Hermitian, so it is in general not an observable, but its *de facto* mean value (which is a complex number) can still be estimated. We discuss this in detail in the proof of Lemma 2.

We are interested in the *output fidelity*, defined as

$$F_{\text{output}} := \frac{\sum_i p_i |\beta_i|^2 |\langle\phi_i'|U|\phi_{\text{in}}\rangle|^2}{\sum_i p_i |\beta_i|^2}. \tag{12}$$

This quantifies the fidelity between the state being measured to generate samples from P_{exp} and the ideal state that can be measured to generate samples from P_{ideal}, and thus can be directly related to the TVD between distributions, $\text{TVD}(P_{\text{exp}}, P_{\text{ideal}})$. In Appendix S2, we explicitly relate F_{output} and $\text{TVD}(P_{\text{exp}}, P_{\text{ideal}})$, and find the threshold fidelity 0.915 using the hardness result proved in [32], which gives a criterion for verified quantum advantage.

In Appendix S2, we also derive a lower bound for F_{output} in terms of $\epsilon := 1/4 - \text{Tr}[\rho O_{10}]$, $\epsilon' := 1/2 - p_{\text{samp}}$, and $\epsilon'' := 1 - F_{\text{in}}(\rho)$, as follows.

Theorem 3 (Lower bound on the output fidelity)**.**

$$F_{\text{output}} \geq 1 - 16\epsilon - 3\epsilon'' + \text{h.o.} \tag{13}$$

where h.o. *indicates higher-order terms in* $\epsilon, \epsilon', \epsilon''$.

If the device is close to perfect (which is the scenario we consider here), then $\epsilon, \epsilon'' \ll 1$ and $|\epsilon'| \ll 1$. Hence, the higher-order terms can be safely dropped, as is shown in detail in Appendix S2, and the above bound can be written as

$$F_{\text{out}}(\rho) \geq 16|\text{Tr}[\rho O_{10}]|^2 + 3F_{\text{in}}(\rho) - 6. \tag{14}$$

Using Theorem 3 with threshold fidelity 0.915, we conclude that the measurement outcomes sample from a classically intractable distribution provided $4|\text{Tr}[\rho O_{10}]|^2 \geq 0.988$ and $F_{\text{in}}(\rho) \geq 0.988$.

Observe that the final lower bound does not contain first-order terms in $\epsilon' = 1/2 - p_{\text{samp}}$. However, we still need to estimate p_{samp} to ensure that its value is sufficiently close to $1/2$ that our first-order approximation still holds. Hence, we also require $|1/2 - p_{\text{samp}}| \leq 0.012$.

It is clear from the above theorem that our protocol can also tolerate a small amount of noise in the measurements of the quantum state. To simplify the analysis, in the rest of this section, we make the perfect-measurement assumption: all measurements, whether performed by the prover in the trusted-measurement scheme or by the verifier in the state-transmission scheme, are noiseless. We postpone the discussion of noisy measurements to Appendix S3.

We claim that the number of copies of the history state needed to verify quantumness (i.e., the sample complexity) depends only on the precision and is not related to the system size n. As a result, the prover only needs to perform $O(n)$ trusted single-qubit measurements. These properties are formalized and proven in Lemmas 1 and 2.

Since the TVD between ideal and real output distributions is lower bounded by estimating F_{in} and $\text{Tr}[\rho O_{10}]$, the sample complexity of the protocol is determined by how many copies of the state are required to estimate both quantities to a specific precision.

Lemma 1 (Sufficiency of single-qubit measurements for F_{in} and p_{samp})**.** *A verifier capable of single-qubit measurements and polynomial-time classical computation can estimate* F_{in} *and* p_{samp} *in a mixed state* ρ *with error at most* δ_o *using* $O(1/\delta_o^2)$ *samples of* ρ.

Proof. First recall that the ideal input state $|\phi_{\rm in}\rangle$ is a product state of either $|x\rangle := \frac{1}{2}\left[(1+\mathrm{i})\,|0\rangle + (1-\mathrm{i})\,|1\rangle\right]$ or $|y\rangle := \frac{1}{2}\left[(1+\mathrm{i})\,|0\rangle + e^{-\mathrm{i}\pi/4}(1-\mathrm{i})\,|1\rangle\right]$. Their corresponding orthogonal states are $\left|x^{\perp}\right\rangle := \frac{1}{2}[(1-\mathrm{i})\,|0\rangle - (1+\mathrm{i})\,|1\rangle]$ and $\left|y^{\perp}\right\rangle := \frac{1}{2}[(1-\mathrm{i})\,|0\rangle - e^{-\mathrm{i}\pi/4}(1+\mathrm{i})\,|1\rangle]$, respectively.

If a pure state $|\psi_i\rangle = \alpha_i\,|0\rangle\,|\phi_i\rangle + \beta_i\,|1\rangle\,|\phi_i'\rangle$ is given, the fidelity of the input state, $|\langle\phi_i|\phi_{\rm in}\rangle|^2$, can be estimated as follows. We first measure the clock qubit in the Z basis, and if the outcome is $+1$ (so the state collapses to $|0\rangle\,|\phi_i\rangle$), we measure every other qubit in its corresponding rotated basis, which is either $\{|x\rangle, \left|x^{\perp}\right\rangle\}$ or $\{|y\rangle, \left|y^{\perp}\right\rangle\}$. If all measurement outcomes are $+1$, then $|\phi_i'\rangle$ collapses to $|\phi_{\rm in}\rangle$. Therefore, if the number of copies for which the clock qubit measurement gives $+1$ is $N_{\rm in+}$, and among them the number of copies where all other measurements give $+1$ is $N_{\rm in+0}$, then $\frac{N_{\rm in+0}}{N_{\rm in+}}$ is an unbiased estimator of $|\langle\phi_i|\phi_{\rm in}\rangle|^2$. Furthermore, for a mixed state ρ, the same strategy gives an estimate of $F_{\rm in}(\rho)$:

$$F_{\rm in}(\rho) = \lim_{N_{\rm in+}\to\infty} \frac{N_{\rm in+0}}{N_{\rm in+}}. \tag{15}$$

The precision of estimating $F_{\rm in}$ increases with $N_{\rm in+}$. More precisely, we can use Hoeffding's inequality to quantify their relationship:

$$\Pr[|F_{\rm in,M} - F_{\rm in}| \geq \delta_o] \leq 2\exp\left(-2\delta_o^2 N_{\rm in+}\right), \tag{16}$$

where $F_{\rm in,M}$ represents the estimate from measurements. For the estimate of $F_{\rm in}$ to have error at most δ_o with probability at least $1-p_e$, it suffices to use $N_{\rm in+} = O(|\ln p_e|/\delta_o^2)$ valid measurements, independent of the system size. Moreover, since the single-step history state has equal weight between the $|0\rangle$ and $|1\rangle$ states of the clock qubit, $N_{\rm in+}$ should be close to $N_{\rm M}/2$, where $N_{\rm M}$ is the total number of states measured.

We also describe how to estimate $p_{\rm samp}$. Fortunately, this can already be obtained from $N_{\rm in+}$. Since $p_{\rm samp}$ is just the probability of a Z-basis measurement of the first qubit returning -1, $\frac{N_{\rm in+}}{N_M}$ is an unbiased estimator of $p_{\rm samp}$. Similarly, the probability for the estimate of $p_{\rm samp}$ to have error more than δ_o is upper bounded as

$$\Pr[|p_{\rm samp,M} - p_{\rm samp}| \geq \delta_o] \leq 2\exp\left(-2\delta_o^2 N_{\rm M}\right), \tag{17}$$

where $p_{\rm samp,M}$ denotes the estimated value of $p_{\rm samp}$. Since $N_{\rm M} > N_{\rm in+}$, we can always estimate $p_{\rm samp}$ to a higher precision than $F_{\rm in}$ when they are estimated together.

□

Lemma 2 (Sufficiency of single-qubit Pauli measurements for $|\langle O_{10}\rangle|^2$)**.** *A verifier capable of single-qubit measurements and polynomial-time classical computation can estimate $|\langle O_{10}\rangle|^2$ in a mixed state ρ with error at most δ_o using $O(1/\delta_o^2)$ samples of ρ.*

Proof. We can write

$$O_{10} = |1\rangle\langle 0| \otimes U = \frac{1}{2}(X - \mathrm{i}Y)\otimes U. \tag{18}$$

It can be difficult to measure O_{10} in general, because U typically decomposes into exponentially many Pauli terms. Fortunately, in our protocol, we have $U = \exp(-\mathrm{i}HT)$ for the ZZ-type Hamiltonian

$$H = \frac{\pi}{4}\sum_{k=1}^{m} H_k = \frac{\pi}{4}\sum_{\{i,j\}\in \rm NN} Z_i Z_j, \tag{19}$$

where each H_k is one of the Z_iZ_js. As all H_k terms commute, we can decompose U into a product of evolutions for each term, and further express these evolutions in terms of trigonometric functions as every H_k is a Pauli string:

$$\begin{aligned} U &= \exp\left(-\mathrm{i}\frac{\pi}{4}\sum_{k=1}^{m} H_k\right) \\ &= \prod_{k=1}^{m} \exp\left(-\mathrm{i}\frac{\pi}{4}H_k\right) \\ &= \prod_{k=1}^{m} \left(\cos\left(\frac{\pi}{4}\right)I - \mathrm{i}\sin\left(\frac{\pi}{4}\right)H_k\right). \end{aligned} \tag{20}$$

U is not a well-defined quantum observable since it is not Hermitian, but we can still define its *de facto* single-measurement outcome as a complex number u. Since all H_ks can be simultaneously measured, u can be inferred by evaluating the right-hand side of Eq. (20). More specifically, letting h_k denote the outcome of a single measurement of H_k, we have

$$u = \prod_{k=1}^{m} \left(\cos\left(\frac{\pi}{4}\right) - \mathrm{i}\sin\left(\frac{\pi}{4}\right)h_k\right). \tag{21}$$

Since each H_k is Z_iZ_j, the verifier need only perform single-qubit Z measurements to obtain the h_ks.

In summary, to estimate the expected value of O_{10}, it suffices to measure the clock qubit in either the X or the Y basis, measure all other qubits in the Z basis to get the values of u, and repeat this process enough times to obtain the mean values of $X\otimes U$ and $Y\otimes U$ with sufficiently high precision.

To determine the number of samples required, we evaluate the probability that the measured value deviates from the expected value using concentration bounds. Note that O_{10} is not Hermitian, so its *de facto* measurement outcomes are complex numbers. Recall that $O_{10} = \frac{1}{2}(X\otimes U - \mathrm{i}Y\otimes U)$, so one sample of the value of O_{10} can be obtained by measuring two copies of the state of interest, and both the real and imaginary parts of the measurement outcome of O_{10} are at most $1/2$. Therefore, for any $0 < \delta_o < 1/2$, letting $\langle\cdot\rangle_{\rm M}$ be the average of

the measurement outcomes after running the experiment N_{M} times, and using Hoeffding's inequality,

$$\begin{aligned}
&\Pr\left[\left||\langle O_{10}\rangle_{\mathrm{M}}|^2 - |\langle O_{10}\rangle|^2\right| \geq \delta_o^2\right] \\
&\quad\leq \Pr[|\operatorname{Re}[\langle O_{10}\rangle_{\mathrm{M}}] - \operatorname{Re}[\langle O_{10}\rangle]| \geq \delta_o] \\
&\qquad + \Pr[|\operatorname{Im}[\langle O_{10}\rangle_{\mathrm{M}}] - \operatorname{Im}[\langle O_{10}\rangle]| \geq \delta_o] \\
&\quad\leq 4\exp\left(-2\delta_o^2 N_{\mathrm{M}}\right).
\end{aligned} \tag{22}$$

In conclusion, to ensure that the error in the estimation of $|\langle O_{10}\rangle|^2$ is less than δ_o with probability at least $1-p_e$, it suffices to measure O_{10} on $2N_{\mathrm{M}} = O\left(|\ln p_e|/\delta_o^2\right)$ copies of the state, irrespective of the size of the system. Moreover, if p_e is a negligible function of the security parameter λ, then N_{M} only needs to scale linearly with λ. In other words, the probability of obtaining a wrong estimate of $|\langle \mathrm{Tr}[\rho O_{10}]\rangle|^2$ converges to 0 exponentially with respect to the number of copies, N_{M}. □

C. Our Protocol

In this subsection, we outline the behavior of the verifier and the prover in our protocol, and present the soundness and completeness conditions.

The verifier first provides the prover with descriptions of H and $|\phi_{\mathrm{in}}\rangle$, and the desired number of copies of the history state N_{M} (whose value is determined in Theorems 4 and 6).

The verifier asks the prover to perform measurements to estimate (or measures by herself if state transmission is allowed) $|\langle O_{10}\rangle|^2$, $N_{\mathrm{in}+0}$, and $N_{\mathrm{in}+}$ from the N_{M} samples to verify the correctness of the output state. She also asks the prover to generate samples by measuring the $|\phi'\rangle$ state conditioned on obtaining -1 from measuring the clock state. Therefore, the verifier should generate two random bits for every state before measuring it.

The first bit, b_{sampling}, determines whether the verifier should ask the prover to generate samples ($b_{\mathrm{sampling}} = 1$) or verify the output state ($b_{\mathrm{sampling}} = 0$). If $b_{\mathrm{sampling}} = 1$, the prover should measure the clock qubit in the standard basis and all system qubits in the Hadamard basis. If the clock is measured to be -1, and if the prover passes the verification protocol, then the outcomes of Hadamard measurements on system qubits are samples from the desired distribution.

When $b_{\mathrm{sampling}} = 0$, the verifier must decide whether to use this copy to estimate $|\langle O_{10}\rangle|^2$ or $F_{\mathrm{in}}(\rho)$ and p_{samp} by generating the other random bit b_{testtype}. If the second random bit, b_{testtype}, is 0, then she estimates $F_{\mathrm{in}}(\rho)$ and p_{samp} by asking the prover to measure the clock qubit in the computational basis and all system qubits in their corresponding basis, updating the values of $N_{\mathrm{in}+0}$ and $N_{\mathrm{in}+}$, as in the proof of Lemma 1. For $b_{\mathrm{testtype}} = 1$, she estimates $|\langle O_{10}\rangle|^2$, so the prover should use the same strategy as in the proof of Lemma 2 to measure the value of U and, subsequently, the values of $X \otimes U$ or $Y \otimes U$.

In the end, the verifier estimates the parameters of interest. As in the proofs of Lemmas 1 and 2, we denote the estimated values of $|\langle O_{10}\rangle|^2$, p_{samp}, and F_{in} by $|\langle O_{10}\rangle_{\mathrm{M}}|^2$, $p_{\mathrm{samp,M}}$, and $F_{\mathrm{in,M}}$, respectively. The verifier then decides to accept or not by checking whether the estimated values are within the acceptance ranges, which are $0.994 \leq 4|\langle O_{10}\rangle_{\mathrm{M}}|^2 \leq 1$, $0.994 \leq F_{\mathrm{in,M}}(\rho) \leq 1$, and $0.494 \leq p_{\mathrm{samp,M}} \leq 0.506$. Note that here we choose more stringent values than the quantum advantage criterion in Theorem 3 such that if the fidelity of the output state is slightly below the quantum advantage criterion, the verifier will reject with high probability. This is related to the *soundness* of the protocol, which is discussed in detail in Theorem 6.

We now present the completeness and soundness properties of the protocol. A proof of quantumness is called *complete* if any honest prover with quantum computational ability (which in our case means being able to prepare the required history state $|\psi_{\mathrm{hist}}\rangle$ with tolerable error, as explained in more detail below) is accepted with probability at least 2/3. It is called *sound* if no prover with only classical polynomial-time computational ability can be accepted with probability higher than 1/3.

Before showing the completeness theorem, we observe that any phase error in the clock qubit does not affect the correctness of sampling, which means that a family of history states can be and should be accepted. In fact, one can easily check that $F_{\mathrm{in}}(|\psi_{\mathrm{hist}}(\theta)\rangle\langle\psi_{\mathrm{hist}}(\theta)|) = 1$ and $4|\mathrm{Tr}[|\psi_{\mathrm{hist}}(\theta)\rangle\langle\psi_{\mathrm{hist}}(\theta)| O_{10}]|^2 = 1$ for all $|\psi_{\mathrm{hist}}(\theta)\rangle := \frac{1}{\sqrt{2}}(|0\rangle|\phi_{\mathrm{in}}\rangle + e^{\mathrm{i}\theta}|1\rangle U|\phi_{\mathrm{in}}\rangle)$, where θ can be any real number. This immediately leads to the following completeness result.

Theorem 4 (Completeness)**.** *If the prover constructs $N_M = 3.5 \times 10^6$ copies of $|\psi_{\mathrm{hist}}(\theta)\rangle$ with a fixed value of θ, then the verifier will accept with probability at least* 2/3.

Proof. We can calculate that $F_{\mathrm{in}}(|\psi_{\mathrm{hist}}(\theta)\rangle\langle\psi_{\mathrm{hist}}(\theta)|) = 1$, $4|\mathrm{Tr}[|\psi_{\mathrm{hist}}(\theta)\rangle\langle\psi_{\mathrm{hist}}(\theta)| O_{10}]|^2 = 1$, and $p_{\mathrm{samp}} = 1/2$. Therefore, it suffices to ensure the probabilities that the measurement errors exceed 0.0015 for $|\langle O_{10}\rangle|^2$, and 0.006 for $F_{\mathrm{in,M}}$ and p_{samp}, are all less than 1/3.

Suppose that of N_{M} available samples, $N_{\mathrm{M}}/2$ are used to generate samples, $N_{\mathrm{M}}/4$ are used to estimate $|\langle O_{10}\rangle|^2$, and $N_{\mathrm{M}}/4$ are used to estimate F_{in}. According to Lemmas 1 and 2, and letting $N_{\mathrm{in}+} = N_{\mathrm{M}}/8$, the probability of rejection is at most $\max\{2\exp(-0.006^2 N_{\mathrm{M}}/4), 4\exp(-0.0015^2 N_{\mathrm{M}}/2)\} = 0.08 < 1/3$. □

However, in a real experiment, it is unlikely for a device to only make one specific error—a phase error on the clock qubit—and to otherwise produce $|\psi_{\mathrm{hist}}(\theta)\rangle$ perfectly. Instead, every experimental platform might have its own pattern of noise with multiple types of errors. Our verification scheme also has some robustness against these more general errors. Here we characterize the robustness for the case where the device can prepare a noiseless initial state but its Hamiltonian evolution has some error.

Protocol to demonstrate quantum advantage by analog quantum simulation

Let H be a Hamiltonian to be simulated, and let $|\phi_{\rm in}\rangle$ be the initial state of Hamiltonian evolution.

1. The verifier initializes counters $s_{\{X,U\}}, s_{\{Y,U\}}, N_X, N_Y, N_{\rm in+}, N_{\rm in+0}$ to 0. She sends N_M and classical descriptions of H and $|\phi_{\rm in}\rangle$ to the prover.
2. The prover creates $N_{\rm M}$ copies of the correct history state $|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle\,|\phi_{\rm in}\rangle + e^{i\gamma}\,|1\rangle\,U\,|\phi_{\rm in}\rangle)$, where γ is a fixed arbitrary phase, and (only in the state-transmission scenario) sends them to the verifier.
3. For each state (in the trusted-measurement scenario) to be measured by the prover or (in the state-transmission scenario) to be received by the verifier:
   (a) The verifier generates 2 random bits $b_{\rm sampling}$ and $b_{\rm testtype}$. If $b_{\rm sampling} = 1$, the verifier will obtain a sample from the distribution. If $b_{\rm sampling} = 0$ and $b_{\rm testtype} = 0$, the verifier will check if the input state is correct. If $b_{\rm sampling} = 0$ and $b_{\rm testtype} = 1$, the verifier will check if the Hamiltonian evolution is correct.
   (b) If $b_{\rm sampling} = 1$, the verifier measures (or asks the prover to measure) the first qubit in the standard basis. If the measurement outcome is -1, then she measures all other qubits in the Hadamard basis and stores the measured bit string.
   (c) If $b_{\rm sampling} = 0$ and $b_{\rm testtype} = 0$, the verifier measures (or asks the prover to measure) the first qubit in the standard basis. If the outcome is $+1$:
      i. The verifier updates the counter by $N_{\rm in+} \leftarrow N_{\rm in+} + 1$.
      ii. The verifier measures (or asks the prover to measure) every other qubit in the following basis: if its initial state is supposed to be $|x\rangle$, then measure it in the $\{|x\rangle, |x^{\perp}\rangle\}$ basis; otherwise, measure it in the $\{|y\rangle, |y^{\perp}\rangle\}$ basis.
      iii. If all outcomes are $+1$, she updates the counter as $N_{\rm in+0} \leftarrow N_{\rm in+0} + 1$.
   (d) If $b_{\rm sampling} = 0$ and $b_{\rm testtype} = 1$:
      i. The verifier selects the basis from $\{X, Y\}$ randomly, measures (or asks the prover to measure) the clock qubit in the chosen basis, and stores the measurement outcome in b.
      ii. The verifier measures (or asks the prover to measure) all system qubits in the Z basis. Then, she calculates the values of U according to the proof of Lemma 2, denoted by u.
      iii. If the basis chosen is X, the verifier updates the counters as $N_X \leftarrow N_X + 1$, $s_{\{X,U\}} \leftarrow s_{\{X,U\}} + bu$.
      iv. If the basis chosen is Y, the verifier updates the counters as $N_Y \leftarrow N_Y + 1$, $s_{\{Y,U\}} \leftarrow s_{\{Y,U\}} + bu$.
4. (a) The verifier calculates $h_{X,U} = s_{\{X,U\}}/N_X$ and $h_{Y,U} = s_{\{Y,U\}}/N_Y$. She also calculates $\langle O_{10}\rangle_{\rm M} = h_{X,U} - {\rm i}h_{Y,U}$ and $4|\langle O_{10}\rangle_{\rm M}|^2$.
   (b) The verifier calculates $F_{\rm in,M} = \frac{N_{\rm in+0}}{N_{\rm in+}}$.
5. If $4|\langle O_{10}\rangle_{\rm M}|^2 > 0.988$ and $F_{\rm in,M} > 0.988$, the verifier accepts the interaction and publishes the stored bit strings as the samples from the distribution. Otherwise, she rejects.

Protocol 1. Our protocol for demonstrating quantum advantage.

Theorem 5 (Completeness + Robustness)**.** *If the prover constructs $N_M = 3.5 \times 10^6$ copies of the noisy history state $|\psi_{\rm noisy}\rangle := \frac{1}{\sqrt{2}}(|0\rangle\,|\phi_{\rm in}\rangle + e^{{\rm i}\theta}\,|1\rangle\,|\phi'\rangle)$ where $|\,\langle\phi'|\,U\,|\phi_{\rm in}\rangle\,|^2 = 0.999$, then the verifier will accept the interaction with probability at least $2/3$.*

Proof. We can check that $F_{\rm in}(|\psi_{\rm noisy}\rangle\langle\psi_{\rm noisy}|) = 1$, $p_{\rm samp} = 1/2$, and $4|\langle O_{10}\rangle|^2 = |\langle\phi'|U|\phi_{\rm in}\rangle|^2 = 0.999$. Therefore, it suffices to estimate $4|\langle O_{10}\rangle|^2$ within precision 0.005 and $F_{\rm in}$ and $p_{\rm samp}$ within precision 0.006. This precision can be achieved using N_M copies of the prepared state, which gives success probability $0.73 > 2/3$. □

Next, we establish the soundness condition. Recall that, informally, a quantum advantage protocol is called *sound* if all provers without quantum computational capability are rejected by the verifier with high probability.

Theorem 6 (Soundness)**.** *If the verifier accepts with probability at least $2/3$ with $N_{\rm M} = 3.5 \times 10^6$ copies of the state provided by the prover, then measurements of the state generate samples from a classically intractable distribution.*

Proof. This theorem has almost been proven in Theorem 3, in which $F_{\rm output} \geq 0.915$ is guaranteed if $F_{\rm in} \geq 0.988$, $|p_{\rm samp} - 1/2| \leq 0.012$, and $4|\langle O_{10}\rangle|^2 \geq 0.988$. Also, according to the proof of Theorem 4, with $N_{\rm M}$ samples, the error in the estimation of all parameters is lower than 0.006 with probability at least $2/3$.

Therefore, if the verifier accepts with probability at least $2/3$, which means that $F_{\rm in,M} \geq 0.994$, $|p_{\rm samp} - 1/2| \leq 0.006$, and $4|\langle O_{10}\rangle_{\rm M}|^2 \geq 0.994$ with probability at least $2/3$, then it is immediately clear that $F_{\rm in} \geq 0.988$, $|p_{\rm samp} - 1/2| \leq 0.012$, and $4|\langle O_{10}\rangle|^2 \geq 0.988$, which im-

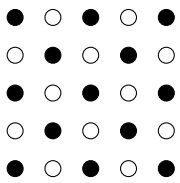

FIG. 2. The square lattice can be divided into two parts such that every ZZ operator acts on qubits from both parts.

plies that $F_{\text{output}} \geq 0.915$. □

A detailed description of the protocol can be found in Protocol 1.

One hidden assumption in this section is that all copies of the history state provided by the prover are independent of each other. However, if the prover is an adversarial challenger, he can provide correlated states. In Appendix S4, we outline how martingale inequalities can be used to show that our protocol is sound even if the states measured are correlated across multiple trials.

The analysis in this section assumes noiseless measurements, which are impractical in real devices. We discuss the protocol's tolerance of noisy measurements in Appendix S3.

III. THE HONEST-PROVER STRATEGY

A. History State Preparation

Our protocol features a rather efficient verification strategy, but for it to be practical, the prover must be able to prepare $O(1)$ copies of the single-step history state of the ZZ-type quantum simulation. A simple approach is to run the time-independent Hamiltonian evolution generated by

$$H_{\text{prep}} = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes H, \tag{23}$$

giving

$$\exp(-\mathrm{i}H_{\text{prep}}T)\left[\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)\,|\phi\rangle\right] = |\psi_{\text{hist}}\rangle\,. \tag{24}$$

However, H_{prep} contains 3-body interaction terms. It is possible for near-term devices to implement a 3-body Hamiltonian (see for example Refs. [39–41]), but it may be challenging to realize H_{prep} in this way.

To circumvent the hardness of implementing 3-body interactions, we propose an echo-based method for preparing history states using only 1-qubit and 2-qubit operations.

One can easily prepare the history state of H by running a half-T evolution of H from the state

$$|\psi'_{\text{hist}}\rangle \propto |0\rangle \exp\left(\tfrac{\mathrm{i}}{2}HT\right)|\phi_{\text{in}}\rangle + |1\rangle \exp\left(-\tfrac{\mathrm{i}}{2}HT\right)|\phi_{\text{in}}\rangle\,. \tag{25}$$

The state $|\psi'_{\text{hist}}\rangle$ can be prepared as follows. Since H involves nearest-neighbor ZZ interactions in a square lattice, one can divide all qubits into two parts such that every ZZ term acts on qubits from different parts, as shown in Fig. 2. Call the filled dots part A, and the non-filled dots part B. Apply CNOT_B gates before and after a $T/2$ time evolution, where CNOT_B is controlled by the clock qubit and acts on the whole part B, followed by an X operation (denoted by X_0) on the clock qubit. This gives the state (up to normalization)

$$\begin{aligned} &X_0\mathrm{CNOT}_B \exp\left(-\tfrac{\mathrm{i}}{2}HT\right)\mathrm{CNOT}_B(|0\rangle + |1\rangle)\,|\phi_{\text{in}}\rangle \\ &= |1\rangle \exp\left(-\tfrac{\mathrm{i}}{2}HT\right)|\phi_{\text{in}}\rangle + |0\rangle\, X_B \exp\left(-\tfrac{\mathrm{i}}{2}HT\right) X_B\,|\phi_{\text{in}}\rangle \\ &= |0\rangle \exp\left(\tfrac{\mathrm{i}}{2}H_2T\right)|\phi_{\text{in}}\rangle + |1\rangle \exp\left(-\tfrac{\mathrm{i}}{2}H_2T\right)|\phi_{\text{in}}\rangle\,, \end{aligned} \tag{26}$$

where X_B denotes X operators acting on all qubits of part B.

One might be concerned that applying CNOT gates on only *half* of the lattice could be difficult with a near-term device. However, one can implement CNOT_B using only a global controlled-Z (CZ) operator and local Hadamard operators H. For all qubits in B, we perform the operation $\mathrm{H}\cdot CZ\cdot\mathrm{H}$, which is exactly a CNOT_B gate. For qubits in A, we do not apply Hadamard operators, so the controlled-Z operation only adds a phase to the second state. This phase is canceled out in the end, because this effective CNOT_B operation is performed twice, and $Z^2 = I$.

Note that this echo approach works for more general Ising-type Hamiltonians, although they might not be easy to verify. A more general discussion can be found in Appendix S5.

In summary, to realize the proposed protocol, the experimental platform should have at least n system qubits and be capable of running single-qubit operations, nearest-neighbor ZZ interactions, and a global CZ operation, which is exactly the capability of our mostly analog + GCZ model of quantum computation. The quantum circuit for a 4-qubit toy model is shown in Fig. 3.

B. Prospects for Experimental Implementation

As explained in Section III A, our protocol uses the mostly analog + GCZ capability, which roughly contains two types of ingredients: first, an analog simulator capable of implementing a ZZ-type Hamiltonian along with a limited number of single-qubit rotations and measurements, and, second, a global CZ gate. The first ingredient is easily accessible in many different hardware platforms including trapped ions, neutral-atom arrays, and superconducting qubits. The second ingredient is not common in hardware architectures for digital quantum computing, but similar ideas have been explored in the context of routing and switching of single- or few-photon signals [42–44] using atomic excitations, and in the case of single-photon-controlled switches [27, 28], where a single photon

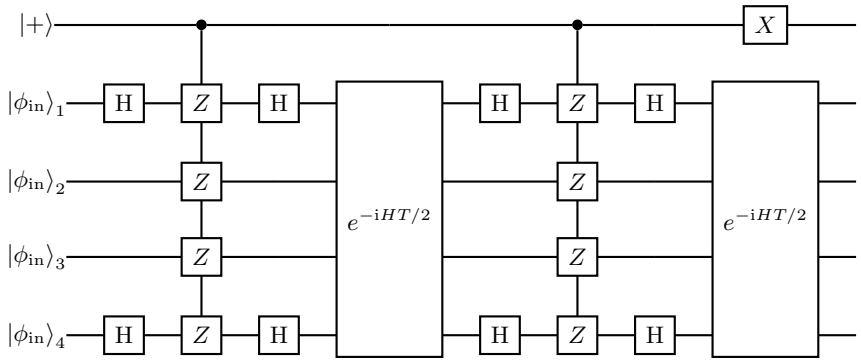

FIG. 3. The final quantum circuit for a (4+1)-qubit example system, where the initial state has been prepared as $|\psi_{\rm initial}\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)\,|\phi_{\rm in}\rangle)$. Here the first qubit is the clock qubit, and part B consists qubits 1 and 4, while part A consists of qubits 2 and 3. The initial state $|\psi_{\rm initial}\rangle$ can be prepared by single-qubit rotations. By applying Hadamard gates before and after the globally controlled-$ZZZZ$ gate for qubits in part B, a controlled-$XZZX$ is implemented. As single-qubit Z commutes with $e^{-{\rm i}HT}$, the Z operations cancel out for qubits in block A.

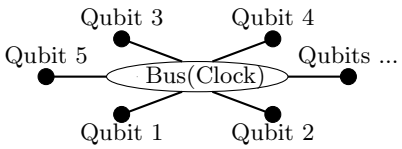

FIG. 4. The "bus" scheme for realizing a global CZ gate. All simulation qubits are only coupled with the central "bus" cavity mode, which behaves effectively as the clock qubit. Both the global CZ gate and the $ZZ+Z$ interaction between simulation qubits can be mediated via the bus mode.

can be used to switch the state of all the photons in a wave packet.

There are two possible ways for the clock qubit to globally turn simulator-qubit Z gates on and off. First, coupling between the clock qubit and the simulator qubits can directly implement the Z gates. Second, the clock qubit can be used to switch classical controls to the simulator qubits on and off. Both implementations are in principle possible, and each of them comes with its own unique set of hardware constraints and challenges, which we describe below.

In the first method, the clock qubit itself is the signal that gives rise to Z rotations of the simulator qubits. This requires the clock qubit to interact directly with *all* of the simulator qubits. Therefore, this method requires the existence of a single global "bus" degree of freedom (e.g., a qubit or bosonic mode) that contains the two clock-qubit states $|0\rangle$ to $|1\rangle$, shown schematically in Fig. 4. The bus mode must interact strongly with each of the system qubits so that a single excitation/photon (or few photons) in the bus mode can produce a significant effect. Such interactions are possible if all of the simulator qubits are strongly coupled to a single cavity mode. Furthermore, the bus-system interaction cannot be resonant (i.e., cannot involve direct absorption of excitations in the bus); otherwise a Z gate on N system qubits cannot be achieved without $M \geq N$ excitations of the bus mode. The bus-system interaction must, instead, be dispersive so that the presence of excitations in the bus mode gives rise to phase shifts of simulator qubits.

While atom-cavity interactions in the single-photon strong-coupling regime are possible in atomic cavity QED, coupling strengths in the so-called strong dispersive regime which are strong enough to produce an off-resonant CZ gate with only a few photons are typically only achievable with superconducting qubits coupled to microwave cavities and superconducting qubits [45]. Alternative implementations may also exist using a confined phonon mode, such as the vibrational modes of an ion trap [46]. Lastly, making use of strong dispersive couplings to implement a globally controlled Z gate between a single excitation in the bus cavity and all of the system qubits would require building the entire simulator inside or attached to the single bus cavity. While this is certainly possible in principle, it is a bespoke feature that would need to be incorporated into the simulator as part of its initial design.

Because of the stringent hardware constraints for the "bus" method of implementing the clock qubit, it is worth considering other methods in which the operation of the clock qubit is more separated from the operation of the quantum simulator under test. Separating these two means that instead of being used to switch the simulator qubits directly, the clock qubit must now switch the *control* signals for single-qubit Z gates on and off. This architecture, shown schematically in Fig. 5, provides significantly more separation between the design constraints of the simulator and those of the clock qubit, but it requires the clock qubit to control a very high-performance quantum switch. In particular, it is not sufficient to use a classical switch with an extremely low switching energy provided by the clock qubit; instead, the switch itself must be able to exist in a superposition between on and off. Such a superposition switching state is extremely

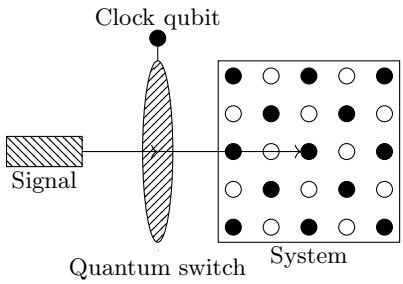

FIG. 5. The quantum switch scheme. Here the simulation qubits are assigned in the square lattice as usual. A photon source gives signals that implement Z operations for each simulation qubit. A high-performance quantum switch, controlled by the clock qubit which could be in superposition, determines whether the signal can be received by simulation qubits or not, which realizes a global CZ gate.

challenging to achieve with large control signals since there are many opportunities to lose a photon (and thus destroy the superposition). Compared to other architectures, superconducting qubits typically require very low switching power—as low as a few photons—so they are a likely candidate for implementation of the necessary quantum switch. For example, a broadband and high-dynamic-range switch such as the one demonstrated in Pechal et al. [47] could be converted to use, e.g., a galvanically coupled fluxonium qubit [48] as the switching element. In the optical domain, single-photon controlled switches have been implemented using atomic ensembles [27] and self-assembled semiconductor quantum dots [28] as the switching medium.

IV. SUMMARY & DISCUSSIONS

In summary, we have proposed a novel scheme for demonstrating quantum computational ability based on verification of analog quantum simulation. The verifier in the scheme need only be capable of polynomial-time classical computation. The prover can be an analog quantum simulator with the additional power of single-qubit operations and a specific global CZ gate, and only needs to be able to prepare a constant number of samples, independent of the system size. Additionally, we assume the prover can perform trusted measurements. We also described some possible near-term experimental implementations of the global CZ gate.

Hangleiter et al. [23] propose another certification scheme that was applied in [24] to verify measurement outcomes using only local measurements. The method in [23] can even verify BQP-complete computation encoded through the Feynman-Kitaev mapping, but it requires $O(n^2)$ samples of the output state for the $ZZ+Z$ Hamiltonian evolution, which is more expensive than our constant-sample-complexity scheme. Our improvement is achieved by a combination of the single-step Feynman-Kitaev encoding and the commuting nature of the $ZZ+Z$ Hamiltonian (or the ZZ Hamiltonian when single-qubit Zs are absorbed). In fact, our protocol can verify all commuting Hamiltonians with constant sample complexity if entangled multi-qubit measurements are allowed, but it is unclear whether there are also near-term honest-prover strategies in this more general case. We discuss this in more detail in Appendix S5.

It is worth noting that there can be a tradeoff between the verification cost and the difficulty of experimental realization. Our verification protocol presented in Section II does not rely on the condition that in H, the coefficient of every term Z_iZ_j is the same ($\pi/4$), but this uniformity makes it possible to simulate the system using $2^{O(\sqrt{n})}$-time classical computation, so that $n = O(\lambda^2)$. If instead the coefficients are randomly selected, then the above simulation is no longer available, and we can conjecture the classical simulation cost to be $2^{\Omega(n)}$, as in [36]. In this case, the number of qubits, the number of single-qubit measurements, and the classical computational cost can all be reduced to $O(\lambda)$—at the cost of more difficult history state preparation—since non-uniform Hamiltonian evolution is in general more challenging.

As the main technical tool of this work, we studied a simplified single-step Feynman-Kitaev construction and developed a scheme to lower bound the output fidelity F_{output} (and subsequently the TVD between ideal and experimental distributions) using three parameters. In fact, the lower bound holds for any unitary U, but the three parameters may not be efficiently estimatable in general. One might ask if we can simply combine the protocol of Fitzsimons et al. [16] with our single-step construction to verify arbitrary quantum operations, such as non-commuting Hamiltonian evolutions or digital quantum circuits. We do not have a definite answer, but this seems difficult for most hard-to-simulate unitaries because they generally decompose into exponentially many Pauli terms and, unlike $ZZ+Z$ or ZZ Hamiltonian evolution, their *de facto* measurement outcome cannot be efficiently deduced from poly(λ) single-qubit measurements.

Experimental implementation of the protocol would be of significant interest. Although it might be difficult to implement quantum communication in the adversarial scenario, our protocol could be a useful tool for experimentalists to benchmark the quality of their devices, because the quality of initial state preparation and that of Hamiltonian evolution can be estimated separately and precisely. As shown in Appendix S2, if the noise pattern is known to be fully stochastic instead of coherent, the experimentalist only needs to achieve output fidelity 0.708, which is significantly easier than the bound of 0.915 in the fully coherent case.

Finally, our approach may have applications to realizing near-term quantum advantage even in devices capable of digital quantum computation. Reconfigurable atom arrays [49, 50] may be one such system. In these arrays, physical qubits (realized by individual neutral atoms con-

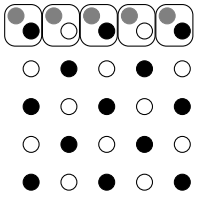

FIG. 6. The reconfigurable atom array scheme. The gray dots represent qubits in a $\sqrt{n}$-qubit GHZ state. Local CZ gates can be realized between pairs of GHZ qubits and system qubits in parallel. Then the GHZ qubits are moved down to the next row and the parallel CZ gates are repeated.

trolled by optical tweezers) can be moved accurately on the 2-D plane in parallel, and transversal CZ gates are available. Therefore, our global CZ gate can be implemented as follows. One can first prepare a large n-qubit GHZ state that behaves as the clock qubit. The GHZ state preparation can be implemented by either performing a sequence of CNOT gates or using constant-depth unitary operations interleaved with measurements and classical computations [51]. One can then can move all qubits in the GHZ state such that every system qubit pairs with a GHZ qubit. Next, using the Levine-Pichler gate [52], CZ gates can be implemented in parallel for every pair of system and GHZ qubits, effectively implementing the global CZ acting on all system qubits. There is also a multi-step solution to mitigate the hardness of GHZ preparation: since our system is a $\sqrt{n} \times \sqrt{n}$ square lattice, it suffices to prepare a 1-D $\sqrt{n}$-qubit GHZ state, and apply the transversal CZ gate $\sqrt{n}$ times to achieve the same global CZ gate. This proposal is depicted in Fig. 6.

While *digital* reconfigurable atom arrays are capable of even more powerful quantum operations than the mostly-analog + GCZ commuting model, it may still be worth performing our proposed experiment using Rydberg atoms. Running our verification protocol gives several quantitative performance measures (F_{in} and $|\langle O_{10}\rangle|^2$), and can thus be used to benchmark the performance of this fast-developing platform in a sample-efficient manner.

ACKNOWLEDGMENTS

We thank Wentai Deng and Ruozhen Gong for helpful discussions. This work received support from the National Science Foundation (QLCI grant OMA-2120757). Z.L. acknowledges financial support by the QuICS Lanczos Graduate Fellowship. D.D. acknowledges support by the NSF GRFP under Grant No. DGE-1840340, an LPS Quantum Graduate Fellowship, and the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Testbed Pathfinder program (award number DE-SC0019040). A.V.G. was also supported in part by AFOSR, DoE ASCR Accelerated Research in Quantum Computing program (award No. DE-SC0020312), DoE ASCR Quantum Testbed Pathfinder program (awards No. DE-SC0019040 and No. DE-SC0024220), NSF STAQ program, AFOSR MURI, and DARPA SAVaNT ADVENT. Support is also acknowledged from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator. D.H. acknowledges funding from the US department of defense through a QuICS Hartree fellowship. Y.K.L. acknowledges support from NIST, and from AFOSR MURI Scalable Certification of Quantum Computing Devices and Networks.

[1] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6): 467–488, 1982.

[2] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.

[3] Seth Lloyd. Universal quantum simulators. *Science*, 273 (5278):1073–1078, 1996.

[4] Andrew M Childs, Dmitri Maslov, Yunseong Nam, Neil J Ross, and Yuan Su. Toward the first quantum simulation with quantum speedup. *Proceedings of the National Academy of Sciences*, 115(38):9456–9461, 2018.

[5] Dominic W. Berry, Graeme Ahokas, Richard Cleve, and Barry C. Sanders. Efficient quantum algorithms for simulating sparse Hamiltonians. *Communications in Mathematical Physics*, 270(2):359–371, 2007.

[6] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103:150502, 2009.

[7] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.

[8] John Preskill. Quantum computing and the entanglement frontier. *arXiv preprint arXiv:1203.5813*, 2012.

[9] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. *Journal of the ACM*, 68(5):1–47, 2021.

[10] Iulia M Georgescu, Sahel Ashhab, and Franco Nori. Quantum simulation. *Reviews of Modern Physics*, 86(1): 153, 2014.

[11] Rainer Blatt and Christian F Roos. Quantum simulations with trapped ions. *Nature Physics*, 8(4):277–284, 2012.

[12] Stuart Flannigan, Natalie Pearson, Guang Hao Low, A Buyskikh, Immanuel Bloch, Peter Zoller, Matthias Troyer, and Andrew J Daley. Propagation of errors and

quantitative quantum simulation with quantum advantage. *Quantum Science and Technology*, 7(4):045025, 2022.

[13] Andrew M. Childs, Yuan Su, Minh C. Tran, Nathan Wiebe, and Shuchen Zhu. Theory of Trotter error with commutator scaling. *Physical Review X*, 11(1), 2021.

[14] Toby S Cubitt, Ashley Montanaro, and Stephen Piddock. Universal quantum Hamiltonians. *Proceedings of the National Academy of Sciences*, 115(38):9497–9502, 2018.

[15] Sepehr Ebadi, Tout T Wang, Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Dolev Bluvstein, Rhine Samajdar, Hannes Pichler, Wen Wei Ho, et al. Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature*, 595(7866):227–232, 2021.

[16] Joseph F Fitzsimons, Michal Hajdušek, and Tomoyuki Morimae. Post hoc verification of quantum computation. *Physical Review Letters*, 120(4):040501, 2018.

[17] Urmila Mahadev. Classical verification of quantum computations. In *Proceedings of the 59th Annual Symposium on Foundations of Computer Science*, pages 259–267. IEEE, 2018.

[18] Alexei Yu Kitaev, Alexander Shen, Mikhail N Vyalyi, and Mikhail N Vyalyi. *Classical and Quantum Computation*. Number 47 in Graduate Studies in Mathematics. American Mathematical Soc., 2002.

[19] Andreas Elben, Benoît Vermersch, Rick van Bijnen, Christian Kokail, Tiff Brydges, Christine Maier, Manoj K. Joshi, Rainer Blatt, Christian F. Roos, and Peter Zoller. Cross-platform verification of intermediate scale quantum devices. *Physical Review Letters*, 124: 010504, 2020.

[20] Jose Carrasco, Andreas Elben, Christian Kokail, Barbara Kraus, and Peter Zoller. Theoretical and experimental perspectives of quantum verification. *PRX Quantum*, 2: 010102, 2021.

[21] J. Ignacio Cirac and Peter Zoller. Goals and opportunities in quantum simulation. *Nature Physics*, 8(44): 264–266, 2012.

[22] Ryan Shaffer, Eli Megidish, Joseph Broz, Wei-Ting Chen, and Hartmut Häffner. Practical verification protocols for analog quantum simulators. *npj Quantum Information*, 7(11):1–12, 2021.

[23] Dominik Hangleiter, Martin Kliesch, Matthias Schwarz, and Jens Eisert. Direct certification of a class of quantum simulations. *Quantum Science and Technology*, 2(1): 015004, 2017.

[24] Juan Bermejo-Vega, Dominik Hangleiter, Martin Schwarz, Robert Raussendorf, and Jens Eisert. Architectures for quantum simulation showing a quantum speedup. *Physical Review X*, 8(2):021010, 2018.

[25] Dorit Aharonov, Michael Ben-Or, Elad Eban, and Urmila Mahadev. Interactive proofs for quantum computations. *arXiv preprint arXiv:1704.04487*, 2017.

[26] Daniel Nagaj. Universal two-body-Hamiltonian quantum computing. *Physical Review A*, 85(3):032330, 2012.

[27] Wenlan Chen, Kristin M Beck, Robert Bücker, Michael Gullans, Mikhail D Lukin, Haruka Tanji-Suzuki, and Vladan Vuletić. All-optical switch and transistor gated by one stored photon. *Science*, 341(6147):768–770, 2013.

[28] Shuo Sun, Hyochul Kim, Zhouchen Luo, Glenn S Solomon, and Edo Waks. A single-photon switch and transistor enabled by a solid-state quantum memory. *Science*, 361(6397):57–60, 2018.

[29] Dan Shepherd and Michael J Bremner. Temporally unstructured quantum computation. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465(2105):1413–1439, 2009.

[30] Michael J Bremner, Richard Jozsa, and Dan J Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2011.

[31] Michael J Bremner, Ashley Montanaro, and Dan J Shepherd. Average-case complexity versus approximate simulation of commuting quantum computations. *Physical Review Letters*, 117(8):080501, 2016.

[32] Martin Ringbauer, Marcel Hinsche, Thomas Feldker, Paul K Faehrmann, Juani Bermejo-Vega, Claire Edmunds, Lukas Postler, Roman Stricker, Christian D Marciniak, Michael Meth, et al. Verifiable measurement-based quantum random sampling with trapped ions. *arXiv preprint arXiv:2307.14424*, 2023.

[33] Gregory D Kahanamoku-Meyer, Soonwon Choi, Umesh V Vazirani, and Norman Y Yao. Classically verifiable quantum advantage from a computational bell test. *Nature Physics*, 18(8):918–924, 2022.

[34] Qingling Zhu, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, Ming Gong, et al. Quantum computational advantage via 60-qubit 24-cycle random circuit sampling. *Science Bulletin*, 67(3):240–245, 2022.

[35] Sergio Boixo, Sergei V Isakov, Vadim N Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J Bremner, John M Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, 2018.

[36] Alexander M Dalzell, Aram W Harrow, Dax Enshan Koh, and Rolando L La Placa. How many qubits are needed for quantum computational supremacy? *Quantum*, 4: 264, 2020.

[37] Han-Sen Zhong, Hui Wang, Yu-Hao Deng, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Jian Qin, Dian Wu, Xing Ding, Yi Hu, et al. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463, 2020.

[38] Yulin Wu, Wan-Su Bao, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, et al. Strong quantum computational advantage using a superconducting quantum processor. *Physical Review Letters*, 127(18):180501, 2021.

[39] HP Büchler, A Micheli, and P Zoller. Three-body interactions with cold polar molecules. *Nature Physics*, 3(10): 726–731, 2007.

[40] Tim Menke, William P Banner, Thomas R Bergamaschi, Agustin Di Paolo, Antti Vepsäläinen, Steven J Weber, Roni Winik, Alexander Melville, Bethany M Niedzielski, Danna Rosenberg, et al. Demonstration of tunable three-body interactions between superconducting qubits. *arXiv preprint arXiv:2205.04542*, 2022.

[41] Bárbara Andrade, Zohreh Davoudi, Tobias Graß, Mohammad Hafezi, Guido Pagano, and Alireza Seif. Engineering an effective three-spin Hamiltonian in trapped-ion systems for applications in quantum simulation. *Quantum Science and Technology*, 7(3):034001, 2022.

[42] Callum R Murray, Alexey V Gorshkov, and Thomas Pohl. Many-body decoherence dynamics and optimized operation of a single-photon switch. *New Journal of Physics*, 18(9):092001, 2016.

[43] Simon Baur, Daniel Tiarks, Gerhard Rempe, and Stephan Dürr. Single-photon switch based on rydberg blockade. *Physical Review Letters*, 112(7):073901, 2014.

[44] Weibin Li and Igor Lesanovsky. Coherence in a cold-atom photon switch. *Physical Review A*, 92(4):043828, 2015.

[45] Alexandre Blais, Arne L Grimsmo, Steven M Girvin, and Andreas Wallraff. Circuit quantum electrodynamics. *Reviews of Modern Physics*, 93(2):025005, 2021.

[46] Christopher Monroe, Wes C Campbell, L-M Duan, Z-X Gong, Alexey V Gorshkov, Paul W Hess, Rajibul Islam, Kihwan Kim, Norbert M Linke, Guido Pagano, et al. Programmable quantum simulations of spin systems with trapped ions. *Reviews of Modern Physics*, 93(2):025001, 2021.

[47] M Pechal, J-C Besse, Mintu Mondal, M Oppliger, S Gasparinetti, and A Wallraff. Superconducting switch for fast on-chip routing of quantum microwave fields. *Physical Review Applied*, 6(2):024009, 2016.

[48] Vladimir E Manucharyan, Jens Koch, Leonid I Glazman, and Michel H Devoret. Fluxonium: Single cooper-pair circuit free of charge offsets. *Science*, 326(5949):113–116, 2009.

[49] Jerome Beugnon, Charles Tuchendler, Harold Marion, Alpha Gaetan, Yevhen Miroshnychenko, Yvan R. P. Sortais, Andrew M. Lance, Matthew P. A. Jones, Gaetan Messin, Antoine Browaeys, and Philippe Grangier. Two-dimensional transport and transfer of a single atomic qubit in optical tweezers. *Nature Physics*, 3(10):696–699, 2007.

[50] Dolev Bluvstein, Simon J Evered, Alexandra A Geim, Sophie H Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, et al. Logical quantum processor based on reconfigurable atom arrays. *Nature*, pages 1–3, 2023.

[51] Adam Bene Watts, Robin Kothari, Luke Schaeffer, and Avishay Tal. Exponential separation between shallow quantum circuits and unbounded fan-in shallow classical circuits. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing*, pages 515–526, 2019.

[52] Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Tout T Wang, Sepehr Ebadi, Hannes Bernien, Markus Greiner, Vladan Vuletić, Hannes Pichler, et al. Parallel implementation of high-fidelity multiqubit gates with neutral atoms. *Physical Review Letters*, 123(17):170503, 2019.

[53] Andreas Elben, Steven T Flammia, Hsin-Yuan Huang, Richard Kueng, John Preskill, Benoît Vermersch, and Peter Zoller. The randomized measurement toolbox. *Nature Reviews Physics*, 5(1):9–24, 2023.

[54] Devdatt P Dubhashi and Alessandro Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.

APPENDICES

In these appendices, we present details omitted from the main text. In Appendix S1, we state and briefly explain the conjectures used to establish computational hardness [24, 32]. In Appendix S2, we lower bound the output fidelity and the total variation distance between distributions using the parameters in our verification scheme. In Appendix S3, we discuss noisy measurements and estimate the noise rate that both verification and sampling can tolerate. In Appendix S4, we discuss an additional soundness property of our protocol against correlated output states using martingale inequalities. In Appendix S5, we generalize the echo method presented in the main text to more general Ising-type Hamiltonians.

Appendix S1: Conjectures for the classical hardness

As mentioned in the main text, the classical hardness of X-basis sampling from a state produced by $(ZZ + Z)$-Hamiltonian evolution is proven in Bermejo-Vega et al. [24] and Ringbauer et al. [32] under several plausible conjectures, which we review in this section.

Conjecture 1 (Polynomial Hierarchy—Conjecture 1 in [24])**.** *The polynomial hierarchy is infinite.*

The second conjecture considers the hardness of a random nearest-neighbor Ising model on an $n \times m$ square lattice where m grows at least linearly with n, with the Hamiltonian

$$H^{(\alpha,\beta)} = \sum_{i,j} \frac{\pi}{4} Z_i Z_j - \sum_i h_i^{(\alpha,\beta)} Z_i, \tag{S1.1}$$

where $h_i^{(\alpha,\beta)} = \frac{\pi}{4} - \frac{\alpha_i+\beta_i}{2}$ with $\alpha_i \in \{0,\pi\}, \beta_i \in \{0,\pi/4\}$ chosen uniformly at random.

Conjecture 2 (Average-case complexity—Conjecture 2 in [24] and conjecture in [32])**.** *Let* $Z^{(\alpha,\beta)} := \mathrm{Tr}\left(e^{\mathrm{i}H^{(\alpha,\beta)}}\right)$. *Approximating* $|Z^{(\alpha,\beta)}|^2$ *up to relative error* $1/4 + o(1)$ *for any 0.001 fraction of the field configurations is #P-hard.*

The last conjecture is about anti-concentration of the output distribution. Consider a one-dimensional nearest-neighbor n-qubit $\Theta(n)$-depth random circuit

$$\mathcal{C} = \left[\prod_{i=1}^{n-1} CZ_{i,i+1}\right]\left[\prod_{i=1}^{n} Z_i^{c_i} e^{-\mathrm{i}\frac{\pi}{4} d_i Z_i} H_i\right], \tag{S1.2}$$

where c_i, d_i are uniformly randomly chosen from $\{0,1\}$ and H_i are Hadamard gates.

Conjecture 3 (Anti-concentration—Conjecture 3 in [24])**.** *For the random circuit* $\mathcal{C}$ *described above,*

$$\Pr_{\mathcal{C}}\left(|\langle x|\mathcal{C}|0\rangle^{\otimes n}|^2 \geq \frac{1}{2^n}\right) \geq \frac{1}{e} \tag{S1.3}$$

for any binary string $x \in \{0,1\}^n$.

Appendix S2: Relating the parameters to the total variation distance

In this appendix, we derive an upper bound on the total variation distance of interest, $\mathrm{TVD}(P_{\mathrm{ideal}}, P_{\mathrm{real}})$, in terms of the parameters F_{in}, $|\langle O_{10}\rangle|^2$, and p_{samp}. We use the same definition of ρ and $|\psi_i\rangle$ as in Eqs. (6) and (7).

First, we relate the TVD and the output fidelity

$$F_{\mathrm{out}}(\rho) := \frac{\sum_i p_i |\beta_i|^2 \left|\langle \phi_i'|U|\phi_{\mathrm{in}}\rangle\right|^2}{\sum_i p_i |\beta_i|^2}. \tag{S2.1}$$

This is the fidelity between the state used for sampling, ρ, and $U|\phi_{\mathrm{in}}\rangle$, since the state corresponding to the "output" of the computation is $|\phi_i'\rangle$ for all i.

In the second step, we derive a lower bound on the state fidelity in terms of the parameters. We lower bound $F_{\mathrm{out}}(\rho)$ using only the parameters $\mathrm{Tr}[\rho O_{10}]$ and $F_{\mathrm{in}}(\rho)$. We find

$$F_{\mathrm{out}}(\rho) \geq 16|\mathrm{Tr}[\rho O_{10}]|^2 + 3F_{\mathrm{in}}(\rho) - 6 \tag{S2.2}$$

up to higher-order terms. As a sanity check, if the history state is perfectly prepared, both $|\mathrm{Tr}[\rho O_{10}]|^2$ and F_{in} should take their maximum values, which are $1/4$ and 1 (as shown later in this section), giving $F_{\mathrm{out}} = 1$ as expected.

1. Relating the total variation distance to the output fidelity

To demonstrate quantum advantage, we generate samples from the desired distribution P_{ideal} defined by $U|\phi_{\text{in}}\rangle$ with total variation distance (TVD) less than $\delta = 0.292$ as per Ringbauer et al. [32]. Therefore, we would like to relate the fidelity F_{out} obtained from the measurements to the distance between the distribution P_{real} corresponding to the classical mixture of $|\phi'_i\rangle$s (i.e., $\sum_i p_i|\beta_i|^2 |\phi'_i\rangle\langle\phi'_i|$) and the ideal distribution P_{ideal}.

Let $\|\cdot\|_{\text{Tr}}$ be the trace norm (Schatten 1-norm). The TVD between probability distributions generated by measurements on quantum states is upper bounded by the trace distance between those states, which is in turn related to the fidelity:

$$\begin{aligned}
\text{TVD}(P_{\text{ideal}}, P_{\text{real}}) &\leq \frac{1}{2}\left\| U|\phi_{\text{in}}\rangle\langle\phi_{\text{in}}|U^\dagger - \sum_i p_i|\beta_i|^2 |\phi'_i\rangle\langle\phi'_i| \right\|_{\text{Tr}} \\
&\leq \sqrt{1 - F\left(U|\phi_{\text{in}}\rangle\langle\phi_{\text{in}}|U^\dagger, \sum_i p_i|\beta_i|^2 |\phi'_i\rangle\langle\phi'_i|\right)} \\
&= \sqrt{1 - \sum_i p_i|\beta_i|^2 F(|\phi'_i\rangle, U|\phi_{\text{in}}\rangle)} = \sqrt{1 - F_{\text{out}}}.
\end{aligned} \tag{S2.3}$$

Thus, $\text{TVD}(P_{\text{ideal}}, P_{\text{real}}) \leq 0.292$ is satisfied if

$$F_{\text{out}} \geq 0.915 > 1 - \delta^2, \tag{S2.4}$$

where we use $\delta = 0.292$.

We also observe that the output fidelity requirement can be relaxed to 0.708 if the noise in the system is known to be fully stochastic. We discuss this in Appendix S2.3.

2. Lower bounding the output fidelity using the parameters

As a mathematical tool, we define an inner product based on the (not explicitly known) diagonalization of ρ. Suppose $\rho = \sum_{i=1}^{2^{n+1}} p_i |\psi_i\rangle\langle\psi_i|$ and there exists an integer $N_{\neq 0} > 1$ such that $p_i > 0$ for all $1 \leq i \leq N_{\neq 0}$ and $p_i = 0$ for all $N_{\neq 0} < i \leq 2^{n+1}$. The inner product $\langle\cdot,\cdot\rangle_\rho$ is defined for the $N_{\neq 0}$-dimensional complex vector space $V = \mathbb{C}^{N_{\neq 0}}$ as

$$\langle \vec{A}, \vec{B}\rangle_\rho := \sum_{1\leq i\leq N_{\neq 0}} p_i A_i B_i^*, \tag{S2.5}$$

where $\vec{A} := (A_1, A_2, \ldots, A_{N_{\neq 0}})^{\text{T}}$ and $\vec{B} := (B_1, B_2, \ldots, B_{N_{\neq 0}})^{\text{T}}$ are vectors in V. It is straightforward to verify that for any valid density matrix ρ, the vector space V equipped with $\langle\cdot,\cdot\rangle_\rho$ is an inner product space. Therefore, one can define the norm of a vector in V as

$$\|\vec{A}\|^2 := \langle \vec{A}, \vec{A}\rangle_\rho = \sum_i p_i |A_i|^2. \tag{S2.6}$$

Next, we define several vectors to help represent the state and the parameters: the *input fidelity vector* $\vec{f}_{\text{in}}$, the *propagation fidelity vector* $\vec{f}_{\text{in}}$, the *output fidelity vector* $\vec{f}_{\text{out}}$, the α *coefficient vector* $\vec{\alpha}$, the β *coefficient vector* $\vec{\beta}$, and the γ *coefficient vector* $\vec{\gamma}$ for a given mixed state ρ, namely

$$\begin{aligned}
\vec{f}_{\text{in}} &:= (\ldots, \langle\phi_i|\phi_{\text{in}}\rangle, \ldots)^{\text{T}}, \\
\vec{f}_{\text{prop}} &:= (\ldots, \langle\phi'_i|U|\phi_i\rangle, \ldots)^{\text{T}}, \\
\vec{f}_{\text{out}} &:= (\ldots, \langle\phi'_i|U|\phi_{\text{in}}\rangle, \ldots)^{\text{T}}, \\
\vec{\alpha} &:= (\ldots, \alpha_i, \ldots)^{\text{T}}, \\
\vec{\beta} &:= (\ldots, \beta, \ldots)^{\text{T}}, \\
\vec{\gamma} &:= (\ldots, \alpha_i\beta_i^*, \ldots)^{\text{T}},
\end{aligned} \tag{S2.7}$$

respectively. Note that $\|\vec{\gamma}\|^2 \leq 1/4$ and $\|\vec{f}_{\text{in}}\|^2, \|\vec{f}_{\text{prop}}\|^2, \|\vec{f}_{\text{out}}\|^2 \leq 1$ since $|\alpha_i|^2 + |\beta_i|^2 = 1$, $\sum_i p_i = 1$, and fidelities are at most 1.

Observe that p_{samp} is the same as $\|\vec{\alpha}\|^2$. Another parameter, $\text{Tr}[\rho O_{10}]$, can be written as the inner product of two of the above vectors:

$$\text{Tr}[\rho O_{10}] = \sum_i p_i \alpha_i \beta_i^* \langle \phi_i' | U | \phi_i \rangle = \langle \vec{\gamma}, \vec{f}_{\text{prop}} \rangle_\rho. \tag{S2.8}$$

Using the Cauchy-Schwarz inequality, we find

$$|\text{Tr}[\rho O_{10}]|^2 = |\langle \vec{\gamma}, \vec{f}_{\text{prop}} \rangle|^2 \leq \|\vec{\gamma}\|^2 \|\vec{f}_{\text{prop}}\|^2 \leq 1/4. \tag{S2.9}$$

Since $\|\vec{\gamma}\|^2 \leq 1/4$ and $\|\vec{f}_{\text{prop}}\|^2 \leq 1$, the above inequality implies that

$$\begin{aligned} 4\|\vec{\gamma}\|^2 &\geq |\text{Tr}[\rho O_{10}]|^2, \\ \|\vec{f}_{\text{prop}}\|^2 &\geq 4|\text{Tr}[\rho O_{10}]|^2. \end{aligned} \tag{S2.10}$$

If the prover performs well, then the estimated value $\text{Tr}[\rho O_{10}]$ should be close to $1/4$, $\|\vec{\alpha}\|^2$ should be close to $1/2$, and F_{in} should be close to 1. Therefore we write them as $\text{Tr}[\rho O_{10}] = 1/4 - \epsilon$, $\|\vec{\alpha}\|^2 = 1/2 + \epsilon' = 1 - \|\vec{\beta}\|^2$, and $F_{\text{in}} = 1 - \epsilon''$, where $\epsilon, \epsilon', \epsilon''$ are all small and $\epsilon, \epsilon'' > 0$. This also implies that $\|\vec{\gamma}\|^2 = \sum_i p_i |\alpha_i|^2 |\beta_i|^2 \geq 1/4 - \epsilon$ and $\|\vec{f}_{\text{prop}}\|^2 \geq 1 - 4\epsilon$.

Recall that our final objective is to lower bound $F_{\text{out}}(\rho)$. We start by giving a lower bound on $\|\vec{f}_{\text{in}}\|^2$ in terms of F_{in}.

First, the Cauchy-Schwarz inequality gives

$$\begin{aligned} F_{\text{in}}(\rho) &= \frac{1}{\|\vec{\alpha}\|^2} \sum_i p_i |\alpha_i|^2 \left| \langle \phi_i | \phi_{\text{in}} \rangle \right|^2 \\ &\leq \frac{1}{\|\vec{\alpha}\|^2} \sum_i p_i |\alpha_i|^2 \left| \langle \phi_i | \phi_{\text{in}} \rangle \right| \\ &\leq \frac{1}{\|\vec{\alpha}\|^2} \left(\sum_i p_i |\alpha_i|^4 \right)^{1/2} \|\vec{f}_{\text{in}}\|. \end{aligned}$$

Plugging in the identity $|\alpha_i|^4 = |\alpha_i|^2 - |\alpha_i|^2 |\beta_i|^2$, we get

$$F_{\text{in}}(\rho) \leq \frac{1}{\|\vec{\alpha}\|^2} (\|\vec{\alpha}\|^2 - \|\vec{\gamma}\|^2)^{1/2} \|\vec{f}_{\text{in}}\|.$$

As before, suppose that $\|\vec{\gamma}\|^2 = 1/4 - \epsilon$ and $\|\vec{\alpha}\|^2 = 1/2 + \epsilon'$. This implies that

$$F_{\text{in}}(\rho) \leq \frac{1}{\frac{1}{2} + \epsilon'} \left(\tfrac{1}{2} + \epsilon' - \tfrac{1}{4} + \epsilon \right)^{1/2} \|\vec{f}_{\text{in}}\|.$$

We can rewrite this as

$$\|\vec{f}_{\text{in}}\| \geq \frac{\frac{1}{2} + \epsilon'}{\sqrt{\frac{1}{4} + \epsilon' + \epsilon}} F_{\text{in}} = (1 - 2\epsilon) F_{\text{in}} + O(\epsilon'^2) + O(\epsilon^2) + O(\epsilon\epsilon').$$

Next, since $|\langle \phi_i' | U | \phi_{\text{in}} \rangle|^2 \geq \left| \langle \phi_i' | U | \phi_i \rangle \right|^2 \left| \langle \phi_i | \phi_{\text{in}} \rangle \right|^2$, we have

$$F_{\text{out}}(\rho) \geq \frac{1}{\|\vec{\beta}\|^2} \sum_i p_i |\beta_i|^2 \left| \langle \phi_i' | U | \phi_i \rangle \right|^2 \left| \langle \phi_i | \phi_{\text{in}} \rangle \right|^2.$$

Note that for any $\delta_1, \delta_2 \in [0, 1]$, we have $(1 - \delta_1)(1 - \delta_2) \geq 1 - \delta_1 - \delta_2 = (1 - \delta_1) + (1 - \delta_2) - 1$. Using this inequality, we can write

$$\begin{aligned} F_{\text{out}}(\rho) &\geq \frac{1}{\|\vec{\beta}\|^2} \sum_i p_i |\beta_i|^2 \left(\left| \langle \phi_i' | U | \phi_i \rangle \right|^2 + \left| \langle \phi_i | \phi_{\text{in}} \rangle \right|^2 - 1 \right) \\ &= -1 + \frac{1}{\|\vec{\beta}\|^2} \sum_i p_i (1 - |\alpha_i|^2) \left(\left| \langle \phi_i' | U | \phi_i \rangle \right|^2 + \left| \langle \phi_i | \phi_{\text{in}} \rangle \right|^2 \right) \\ &= -1 + \frac{1}{\|\vec{\beta}\|^2} \left(\|\vec{f}_{\text{prop}}\|^2 + \|\vec{f}_{\text{in}}\|^2 \right) - \frac{1}{\|\vec{\beta}\|^2} \sum_i p_i |\alpha_i|^2 \left| \langle \phi_i' | U | \phi_i \rangle \right|^2 - \frac{\|\vec{\alpha}\|^2}{\|\vec{\beta}\|^2} F_{\text{in}}(\rho). \end{aligned}$$

The second-to-last term can be bounded in terms of $\|\vec{f}_{\text{prop}}\|$, using the same argument we used to relate $F_{\text{in}}(\rho)$ and $\|\vec{f}_{\text{in}}\|$. This yields

$$\frac{1}{\|\vec{\beta}\|^2}\sum_i p_i|\alpha_i|^2\left|\langle\phi_i'|\,U\,|\phi_i\rangle\right|^2 \leq \frac{\sqrt{\frac{1}{2}+\epsilon'-\frac{1}{4}+\epsilon}}{\frac{1}{2}-\epsilon'}\|\vec{f}_{\text{prop}}\| = (1+4\epsilon'+2\epsilon)\|\vec{f}_{\text{prop}}\| + O(\epsilon'^2)+O(\epsilon^2)+O(\epsilon\epsilon'). \tag{S2.11}$$

Plugging this into the preceding equation, we get

$$\begin{aligned} F_{\text{out}}(\rho) &\geq -1 + \frac{1}{\|\vec{\beta}\|^2}\left(\|\vec{f}_{\text{prop}}\|^2+\|\vec{f}_{\text{in}}\|^2\right) - \frac{\sqrt{\frac{1}{2}+\epsilon'-\frac{1}{4}+\epsilon}}{\frac{1}{2}-\epsilon'}\|\vec{f}_{\text{prop}}\| - \frac{\|\vec{\alpha}\|^2}{\|\vec{\beta}\|^2}F_{\text{in}}(\rho) \\ &\geq -1 + \frac{2}{1-2\epsilon'}\left(\|\vec{f}_{\text{prop}}\|^2+\|\vec{f}_{\text{in}}\|^2\right) - \frac{\sqrt{\frac{1}{2}+\epsilon'-\frac{1}{4}+\epsilon}}{\frac{1}{2}-\epsilon'}\|\vec{f}_{\text{prop}}\| - \frac{\frac{1}{2}+\epsilon'}{\frac{1}{2}-\epsilon'}F_{\text{in}}(\rho) \\ &= 1 - 16\epsilon - 3\epsilon'' + \text{h.o.}, \end{aligned} \tag{S2.12}$$

where h.o. indicates higher-order terms in $\epsilon, \epsilon', \epsilon''$. Numerically, this first-order approximation of the lower bound has absolute error at the 10^{-3} order of magnitude if all of $\epsilon, |\epsilon'|, \epsilon''$ are upper bounded by 0.02. We have thus established Theorem 3.

3. Relaxing the fidelity requirement for fully stochastic noise models

We notice that inequality (S2.3) can be improved to get a bound that approaches

$$\text{TVD}(P_{\text{ideal}}, P_{\text{real}}) \leq 1 - F_{\text{out}} \tag{S2.13}$$

in cases where the errors are stochastic rather than coherent. Let $\rho_{\text{real}} := \sum_i p_i|\beta_i|^2\,|\phi_i'\rangle\langle\phi_i'|$ be the real state (that is, the state prepared in the experiment), and let $\sigma = |\psi\rangle\langle\psi|$ be the ideal pure state. The real state ρ_{real} has fidelity $F_{\text{out}} = \langle\psi|\,\rho_{\text{real}}\,|\psi\rangle = 1-\delta_f$ (where δ_f is the "infidelity").

Furthermore, assume that ρ_{real} is mixed, in the sense that $\text{Tr}\big(\rho_{\text{real}}^2\big) = 1-\delta_p$ (where δ_p is the "impurity"). This assumption can be checked by estimating $\text{Tr}\big(\rho_{\text{real}}^2\big)$ using either randomized measurements [53] or the swap test. (The former method is appropriate for small quantum systems where the experimenter has a relatively limited degree of control; the latter method is capable of handling much larger quantum systems, but requires more sophisticated quantum control.)

Define projectors $\Pi_0 := |\psi\rangle\langle\psi|$ and $\Pi_1 := I - \Pi_0$. Write the state in block-diagonal form as $\rho_{\text{real}} = \rho_{00}+\rho_{01}+\rho_{10}+\rho_{11}$, where $\rho_{ab} := \Pi_a\rho_{\text{real}}\Pi_b$ for $a,b\in\{0,1\}$.

Let $\|\cdot\|_F$ be the Frobenius norm (i.e., the Schatten 2-norm). Then we can upper bound the trace distance between ρ_{real} and σ as follows:

$$\begin{aligned} \|\rho_{\text{real}}-\sigma\|_{\text{Tr}} &\leq \|\rho_{00}-\sigma\|_{\text{Tr}} + \|\rho_{11}\|_{\text{Tr}} + \|\rho_{01}\|_{\text{Tr}} + \|\rho_{10}\|_{\text{Tr}} \\ &= 2\delta_f + 2\|\rho_{01}\|_{\text{Tr}}. \end{aligned} \tag{S2.14}$$

We have

$$\begin{aligned} \|\rho_{01}\|_{\text{Tr}} &= \|\rho_{01}\|_F \\ &= \tfrac{1}{\sqrt{2}}(\text{Tr}\big(\rho_{\text{real}}^2\big) - \|\rho_{00}\|_F^2 - \|\rho_{11}\|_F^2)^{1/2} \\ &\leq \tfrac{1}{\sqrt{2}}(\text{Tr}\big(\rho_{\text{real}}^2\big) - \|\rho_{00}\|_F^2)^{1/2} \\ &= \tfrac{1}{\sqrt{2}}(1-\delta_p-(1-\delta_f)^2)^{1/2} \\ &= \tfrac{1}{\sqrt{2}}(2\delta_f-\delta_f^2-\delta_p)^{1/2}. \end{aligned} \tag{S2.15}$$

Therefore,

$$\frac{1}{2}\|\rho_{\text{real}}-\sigma\|_{\text{Tr}} \leq \delta_f + \sqrt{\delta_f - \delta_f^2/2 - \delta_p/2}. \tag{S2.16}$$

This bound can be compared to inequalities (S2.3) and (S2.13). When ρ is a pure state, we have $\delta_p = 0$, so the above bound is roughly $\sqrt{\delta_f}$, which looks like inequality (S2.3). When ρ is highly mixed, δ_p can be as large as $\delta_p \approx 2\delta_f - \delta_f^2$, so the above bound is roughly δ_f, which looks like inequality (S2.13). This implies that, when the noise model is known to be fully stochastic, the output state fidelity need only be at least $1 - \delta = 0.708$ to demonstrate quantum advantage, according to inequality (S2.13).

Appendix S3: Noisy Measurements

In the analysis in the main article, we assume that all measurements are perfect. In this appendix, we discuss the potential negative effects of noisy measurements in both verification and sampling. We also show that the tolerable noise rate in measurements for an n-qubit system is $\epsilon \ll 1/n$.

1. Noisy measurements in verification

Let us first discuss the estimation of $|\langle O_{10}\rangle|^2 = |\langle X \otimes U\rangle + \mathrm{i}\langle Y \otimes U\rangle|^2$. When $\epsilon \ll 1/n$, the number of erroneous measurements in each estimation of the *de facto* value of $X \otimes U$ or $Y \otimes U$ is much less than 1. Therefore, the mean values measured for both quantities only deviate by up to $n\epsilon\langle X \otimes U\rangle$ and $n\epsilon\langle Y \otimes U\rangle$ due to the measurement errors, leading to constant-factor errors in the estimation of $|\langle O_{10}\rangle|^2$. Hence, the error rate must be sufficiently small, e.g., $\epsilon = \frac{1}{100n}$, such that the estimated value can still be in the range of acceptance.

Similarly, we require the measurement error to be as small as $\frac{1}{100n}$ to estimate F_{in} to sufficiently high precision, because the value of $N_{\mathrm{in}+0}$ could be lowered by $N_{\mathrm{M}} n\epsilon$ when measuring N_{M} samples. This may lead to a constant-factor error (of order $n\epsilon$) in $F_{\mathrm{in,M}}$.

2. Noisy measurements in sampling

In the following lemma, we show that we can still sample from a classically intractable distribution if the measurement error is much lower than $1/n$.

Lemma 3. *If $F_{\mathrm{output}} = 1 - \delta_f$, and all measurements have the same error rate $\epsilon \ll 1/n$, then the measurement outcomes sample from a distribution P_{real} with $\mathrm{TVD}(P_{\mathrm{real}}, P_{\mathrm{ideal}}) \leq \delta' = \sqrt{\delta_f} + O(1)$.*

Proof. Since there are n Hadamard measurements to be performed, the probability of having no error in the measurements is

$$p_{\mathrm{measure}} = (1-\epsilon)^n \approx 1 - \epsilon n. \tag{S3.1}$$

Therefore, there is a $1 - \epsilon n$ probability that the measurement outcome samples from a distribution that is $\sqrt{\delta_f}$ away from the ideal distribution in terms of TVD. In the worst case, we simply assume the distribution of errorneous measurements has maximum TVD from the ideal distribution, which is 1. Hence, the TVD between the real experiment distribution and the ideal distribution can be upper bounded by

$$\mathrm{TVD}(P_{\mathrm{real}}, P_{\mathrm{ideal}}) \leq (1-\epsilon n)\,\mathrm{TVD}(P_{\mathrm{real}}, P_{\mathrm{ideal}}) + \epsilon n = (1-\epsilon n)\sqrt{\delta_f} + \epsilon n = \sqrt{\delta_f} + O(1), \tag{S3.2}$$

where in the last step we use $\epsilon \ll 1/n$ and $\delta_f < 1$. □

Appendix S4: Relaxing the assumption that the trials are i.i.d.

Our protocol consists of N_M repeated trials or experiments that are carried out by the prover and the verifier. In the preceding discussion, we have assumed that these trials are independent and identically distributed (i.i.d.), so that the accuracy of our protocol can be shown using simple large-deviation bounds, such as Hoeffding's inequality. Here we sketch how this i.i.d. assumption can be relaxed. In this case, the accuracy of our protocol can be shown using large-deviation bounds based on martingales, such as Azuma's inequality [54].

To demonstrate this, consider a protocol that estimates the expectation value of an observable A by repeating an experiment (preparing a quantum state and measuring it) N_M times. More complicated protocols can be handled in a similar way. For $j = 1, 2, \ldots, N_M$, let F_j be the random variable that represents the classical measurement outcome

from the jth repetition of the experiment. Let $F = (1/N_M)\sum_{j=1}^{N_M} F_j$ be the average of the F_j, which we use to estimate the expectation value of A. In addition, assume that the operator norm of A is bounded by $\|A\| \leq \beta$, where β is independent of the size of the system, and hence $|F_j| \leq \beta$. This assumption is satisfied for many commonly-used measurements, such as computational-basis measurements preceded by arbitrary single-qubit rotations.

In the case where the trials are i.i.d., the same quantum state ρ is prepared in every trial, and the random variables F_j are i.i.d. with expectation value $\mathrm{Tr}(A\rho)$. Then F has expectation value $\mathrm{Tr}(A\rho)$, and Hoeffding's inequality implies that F satisfies a Gaussian-like tail bound with width $O(\beta/\sqrt{N_M})$.

In the non-i.i.d. case, it is possible for the N_M trials to be correlated. Without loss of generality, we can imagine that there exists a joint state σ on N_M copies of the quantum system, and for each j, the random variable F_j comes from measuring the reduced state on the jth copy of the system, which we denote $\sigma_j := \mathrm{Tr}_{\{1,\ldots,N_M\}\setminus\{j\}}(\sigma)$.

Despite these complications, it is still possible to interpret F as an estimate of the expectation value of A for a particular quantum state τ on a single copy of the system. This follows since each F_j has expectation value $\mathrm{Tr}(A\sigma_j)$, and hence F has expectation value $(1/N_M)\sum_{j=1}^{N_M}\mathrm{Tr}(A\sigma_j) = \mathrm{Tr}(A\tau)$, where $\tau = (1/N_M)\sum_{j=1}^{N_M}\sigma_j$.

Furthermore, despite the fact that the random variables F_j are correlated, one can still show that F satisfies a Gaussian-like tail bound with width $O(\beta/\sqrt{N_M})$. Intuitively, this is because each F_j can influence the value of F by an amount that is bounded by $\pm\beta/N_M$. Formally, this can be shown by well-known martingale techniques (see [54]), i.e., constructing the Doob martingale $G_j = \mathbb{E}(F|F_j,\ldots,F_1)$, showing that G_j has bounded differences $|G_j - G_{j-1}| \leq 2\beta/N_M$, and applying Azuma's inequality.

Appendix S5: Echo for more general Hamiltonians

In the main text, we have shown that the echo approach can be used to generate the single-step history state for a $(ZZ+Z)$-type Hamiltonian on a bipartite lattice. In this section, we show that the single-step history state can be prepared for some—though not all—other Ising-type Hamiltonians.

A $(ZZ+Z)$-type Hamiltonian is very special because its terms commute. This allows us to run the controlled-Zs independently and only worry about controlled-ZZs. For more general non-commuting Hamiltonians, we may have to "invert" all its terms in the echo approach. Under suitable conditions, we can do this using the following theorem.

Theorem 7. *If there exists an operator P which is a product of single-qubit operations such that $PHP = -H$, then the single-step history state can be prepared using 2-local operations and controlled-P gates.*

Proof. We start with the initial state $(|0\rangle + |1\rangle)|\phi\rangle$ and perform CP before and after a half-time evolution of H, followed by a Pauli-X on the clock qubit and a half-time evolution of H. The final state is

$$\begin{aligned}
e^{-\mathrm{i}HT/2}\cdot X_0\cdot CP\cdot e^{-\mathrm{i}HT/2}\cdot CP\,(|0\rangle+|1\rangle)|\phi\rangle &= e^{-\mathrm{i}HT/2}\left[|1\rangle e^{-\mathrm{i}HT/2}|\phi\rangle + |0\rangle Pe^{-\mathrm{i}HT/2}P|\phi\rangle\right]\\
&= e^{-\mathrm{i}HT/2}\left[|1\rangle e^{-\mathrm{i}HT/2}|\phi\rangle + |0\rangle\sum_k \frac{1}{k!}P(-\mathrm{i}HT/2)^k P|\phi\rangle\right]\\
&= e^{-\mathrm{i}HT/2}\left[|1\rangle e^{+\mathrm{i}HT/2}|\phi\rangle + |0\rangle\sum_k \frac{1}{k!}(+\mathrm{i}HT/2)^k|\phi\rangle\right]\\
&= e^{-\mathrm{i}HT/2}\left[|1\rangle e^{-\mathrm{i}HT/2}|\phi\rangle + |0\rangle e^{+\mathrm{i}HT/2}|\phi\rangle\right]\\
&= |0\rangle|\phi\rangle + |1\rangle e^{-\mathrm{i}HT}|\phi\rangle\,,
\end{aligned} \tag{S5.1}$$

which is the desired output. □

There are several cases in which an operator P satisfying the conditions of the theorem can be constructed. For example, if the Hamiltonian consists of ZZ terms on a bipartite interaction graph, then P can apply an X (or Y) operator to all qubits on one half of the bipartition. We can also handle some cases where the Hamiltonian is non-commuting, such as

$$H = \sum_{(i,j)\in \mathrm{NN}} (X_iX_j + Y_iY_j) + \sum_i Z_i \tag{S5.2}$$

acting on a bipartite lattice. Then we can split the system into two sets of qubits where all interactions are between qubits in different sets. If the operator P acts with X on the first set of qubits and Y on the second set, then it anticommutes with each term of H, so it has the desired behavior.