




Preparation and verification of tensor network states

Esther Cruz ^{1,2} Flavio Baccari,^{1,2} Jordi Tura,^{1,2,3} Norbert Schuch ^{1,2,4,5} and J. Ignacio Cirac ^{1,2}¹Max-Planck-Institute of Quantum Optics, Hans-Kopfermann-Straße 1, 85748 Garching, Germany²Munich Center for Quantum Science and Technology, Schellingstraße 4, 80799 München, Germany³Instituut-Lorentz, Universiteit Leiden, P.O. Box 9506, 2300 RA Leiden, The Netherlands⁴University of Vienna, Faculty of Mathematics, Oskar-Morgenstern-Platz 1, 1090 Wien, Austria⁵University of Vienna, Faculty of Physics, Boltzmannngasse 5, 1090 Wien, Austria

(Received 19 August 2021; revised 21 February 2022; accepted 5 May 2022; published 31 May 2022)

We consider a family of tensor network states defined on regular lattices that come with a natural definition of an adiabatic path to prepare them. This family comprises relevant classes of states, such as injective matrix product and projected entangled-pair states, and some corresponding to classical spin models. We show how uniform lower bounds to the gap of the parent Hamiltonian along the adiabatic trajectory can be efficiently computed using semidefinite programming. This allows one to check whether the adiabatic preparation can be performed efficiently with a scalable effort. We also derive a set of observables whose expectation values can be easily determined and that form a complete set, in the sense that they uniquely characterize the state. We identify a subset of those observables which can be efficiently computed if one has access to the quantum state and local measurements, and analyze how they can be used in verification procedures.

DOI: [10.1103/PhysRevResearch.4.023161](https://doi.org/10.1103/PhysRevResearch.4.023161)

I. INTRODUCTION

The simulation of many-body states is one of the most promising and long-awaited applications of quantum computing. In particular, quantum computers are expected to efficiently prepare certain entangled multipartite states, which can help us in the study of quantum many-body systems, or in variational quantum algorithms. The advent of the first generations of both analog and quantum computers has triggered a strong interest in the preparation of such states. For instance, GHZ states up to tens of qubits have been prepared with trapped ions [1–5], Rydberg atoms [6], superconducting qubits [7–10], photons [11–14], or nuclear spins [15,16].

Tensor network states (TNS) constitute an especially appealing family of multipartite states [17]. On the one hand, they are expected to efficiently approximate the ground states of local Hamiltonians. On the other, many paradigmatic states in the realm of quantum information or condensed matter physics are simple examples of TNS. The best-known class of such states is that of matrix product states (MPS) [18], which corresponds to a one-dimensional geometry. Higher-dimensional generalizations are known as projected entangled-pair states (PEPS) [19]. In both cases, they are characterized by the bond dimension D , which is directly related to the size of the tensors building such states. Those states possess a special property, namely, that they are the ground state of a local, frustration-free Hamiltonian. This

implies that, in case of no degeneracy, one can easily check the successful preparation of the state by measuring a set of local observables. Thus, such states naturally play an essential role in the certification of quantum computers [20–32].

The preparation of TNS has been actively pursued in the last years and, in particular, methods that operate efficiently in terms of the number of qudits (or lattice size) N have been devised.¹ Matrix product states can be sequentially generated [33] in a time that scales linearly with N . In fact, MPS of up to 10 qubits have been recently prepared in a superconducting setup [34–37]. Certain kinds of PEPS (the so-called sequentially generated) can also be prepared in the same timescale [38] and proposals for the generation of sequentially generated PEPS have been recently put forward [39,40]. While containing many paradigmatic examples of TNS, those states are fine tuned in the sense that a small change on the tensors defining the state may lead to another PEPS outside that class. In fact, those tensors are strongly restricted by the fact that the states have to be sequentially generated.

In [41], a very efficient quantum algorithm to generate a wide range of PEPS was introduced. That class of states is stable under deformations of the tensors and thus they are not fine tuned. The algorithm is based on the adiabatic method and the circuit depth scales as $O[\ln(N)]$. The algorithm needs, however, the existence of a gap Δ along the adiabatic path, something which is difficult to ensure since checking that typically requires a computational time that scales exponentially with N . Additionally, it is devised for digital quantum computers, but not analog ones.

¹By efficient we mean that the computational time grows at most polynomially with N . We will also use the shortcut “exponential time” meaning that the time scales exponentially with N .

In this paper we consider a family of states on arbitrary lattices and prove two results on this family. First, we show how the computation of the gap of the parent Hamiltonian can be expressed as a semidefinite programming problem (SDP), and how this allows us to efficiently compute lower bounds $\delta \leq \Delta$ on the gap. Second, we show that for such families of states, it is possible to predict the expectation values of many observables beyond those appearing as terms of the parent Hamiltonian. In fact, there is an exponential number of such observables, which forms a complete set in the set of operators acting on the many-body Hilbert space.

The first result naturally leads to preparation protocols by using, for instance, the one in Ref. [41] since we can efficiently identify the subset of tensor network states for which a bound in the gap is known. Besides, we also extend the adiabatic algorithm to continuous time, which is more suitable for analog quantum computers.

The second result naturally leads to certification protocols based on interactive proofs (see [20–26,29–32] for previous works on interactive verification schemes) which are inspired by the difficulty of sampling from PEPS [42,43]. While they require a certification time that grows exponentially with N [31,32], we propose efficient versions that, however, rely on stronger standard complexity assumptions. We also explain that, in case they can be spoofed, this would immediately lead to new classical algorithms to estimate physically relevant expectation values of observables in the class of states we consider.

This paper is structured as follows. In Sec. II we present the class of states and their corresponding parent Hamiltonian. The states depend on two positive parameters, $t, \beta \geq 0$, that can be viewed as time and inverse temperature, respectively. We show that, by construction, these states can be smoothly connected to a product state. This family of states is essentially the same as the one considered in Ref. [41], although our formulation allows to extend the adiabatic quantum algorithm of [41] to continuous time in a natural fashion. We will see later how this family includes injective MPS and PEPS, as well as some ground states of classical models whose PEPS description might not be injective (although their parent Hamiltonian has unique ground state). In Sec. III we first show how one can efficiently find lower bounds on the gap of the parent Hamiltonian by means of an SDP. In particular, for every value of t we can find a maximum value of $\beta(t)$ such that for all $\beta < \beta(t)$ the gap of the Hamiltonian can be lower bounded by a constant that does not depend on the system size. We then introduce sets of operators whose expectation values can be easily computed and that are complete, in the sense that they provide a tomography of the state. Finally, we propose verification protocols in Sec. IV and discuss some possible complexity arguments.

II. STATES AND PARENT HAMILTONIAN

In this section, we introduce the family of states that we will consider in this work, which is comparable to the one considered in Ref. [41]. It is built in terms of sets of local commuting operators, together with two parameters $t, \beta \geq 0$. For $t = \beta = 0$, they are product states, whereas otherwise they are entangled and can be efficiently expressed as PEPS.

Based on that fact, we will explicitly construct a frustration-free parent Hamiltonian, which will play an important role in the procedure to prepare the states. We then analyze how to prepare these family of states adiabatically and study the scaling of the computational time as a function of the system size and a lower bound on the gap. We build upon the work done in Ref. [41] and extend their adiabatic algorithm to continuous times. We review here the main idea of the algorithm from Ref. [41] and its runtime, and refer the reader to the original paper for a more technical discussion. Note that while in Ref. [41] a constant gap was assumed, we will present later in Sec. III a method for lower bounding such gap, which constitutes our main result regarding ground-state preparation. This will immediately allow us to know which states we can efficiently prepare. Lastly, we show how the presented family of states includes many physically relevant examples of TN states.

A. Setting

We will consider a rather general setup, although we will give examples later for regular lattices. We consider N qudits, with Hilbert space $\mathcal{H}_d = \mathbb{C}^d$ and computational basis $\{|0\rangle, \dots, |d-1\rangle\}$, located at the vertices \mathcal{V} of a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with edges \mathcal{E} of bounded degree z (that is, the maximal number of edges starting from a given vertex). We define $\mathcal{H} = \mathcal{H}_d^{\otimes N}$, and denote the set of Hermitian operators acting on \mathcal{H} by \mathcal{A} . For any $O \in \mathcal{A}$, we define its support $\lambda(O) \subset \mathcal{V}$ as the subset of vertices such that $O = \text{tr}_i(O) \otimes \mathbb{1}_i/d$ if and only if $i \notin \lambda(O)$, where tr_i is the trace with respect to the qudit i .

The graph \mathcal{G} defines a natural distance $d(i, j)$ between two vertices $i, j \in \mathcal{V}$ as the minimal number of edges connecting them. We also define the radius of a subset of vertices $\lambda \subset \mathcal{V}$,

$$r(\lambda) = \min_{i \in \lambda} \max_{j \in \lambda} d(i, j). \tag{1}$$

We denote by $\mathcal{A}_r \in \mathcal{A}$ the set of Hermitian operators acting on \mathcal{H} whose support has a radius of at most r .

B. States

We call $\mathbb{K}_{r,M} \subset \mathcal{A}$ the set of operators that can be written as

$$K = \sum_{n=1}^M \kappa_n, \quad [\kappa_n, \kappa_m] = 0, \quad n, m = 1, \dots, M, \tag{2}$$

where $\kappa_n \in \mathcal{A}_r$, $\|\kappa_n\|_\infty \leq 1$, and $\kappa_n \geq 0$. That is, it is the set that can be written as a sum of M commuting, positive-semidefinite, subnormalized operators that have a radius of at most r . Apart from the trivial cases, where the operators κ_n act on single vertices or when they are products of the same single-qudit operator, one can easily construct nontrivial sets $\mathbb{K}_{r,M}$. In Appendix A we briefly review some of them.

Given $r_\alpha, M_\alpha \in \mathbb{N}$, $K_\alpha \in \mathbb{K}_{r_\alpha, M_\alpha}$, for $\alpha = 1, 2$, we define the family of states

$$|\Psi(\beta, t)\rangle = \frac{1}{Z(\beta, t)} e^{\beta K_1} e^{itK_2} |\varphi_1\rangle \otimes \dots \otimes |\varphi_N\rangle, \tag{3}$$

where $\beta, t \geq 0$, Z is a normalization constant, and $|\varphi_i\rangle$ are arbitrary single-qudit states. This family of states obviously




FIG. 1. Example of operators $e^{i\kappa_{2,n}t}$ (straight lines) and $e^{\beta\kappa_{1,m}}$ (wiggled lines) that act each on two adjacent sites. (a) All operators $e^{i\kappa_{2,n}t}$ that act on site j . Here $\mu_j = \{p, q, r, s\}$. (b) Support of the operator $\prod_{n \in \mu_j} e^{-i\kappa_{2,n}t}$. (c) Support of the operator $\prod_{m \in \nu_j} e^{-\beta\kappa_{1,m}}$.

contains all product states if we take $\beta = t = 0$. For $\beta = 0$, we have $Z(0, t) = 1$. Note that we only explicitly denote the dependence of $|\Psi(\beta, t)\rangle$ on β and t , while omitting the dependence on K_α , for $\alpha = 1, 2$, and $|\varphi_i\rangle$, in order to ease the notation.

C. Parent Hamiltonian

We now show that any state (3) is the unique ground state of a local, frustration-free Hamiltonian, which we construct explicitly. We denote by $\kappa_{\alpha,n}$ the operators appearing in the decomposition (2) of K_α , for $\alpha = 1, 2$. For $j \in \mathcal{V}$, let $\mu_j := \{n \mid j \in \lambda(\kappa_{2,n})\}$ be the index set of all terms $\kappa_{2,n}$ which act nontrivially on site j , and $\nu_j = \{m \mid \exists n \in \mu_j : \lambda(\kappa_{1,m}) \cap \lambda(\kappa_{2,n}) \neq \emptyset\}$ the index set of all terms $\kappa_{1,m}$ which overlap with one of the previous terms. See Fig. 1 for an example of how such set of vertices would be. Then, for $j = 1, \dots, N$, we define

$$h_j = O_j^\dagger \Pi_j O_j, \tag{4}$$

where $\Pi_j = \mathbb{1}_j - |\varphi_j\rangle\langle\varphi_j|_j$ acts on the qudit at vertex j , and

$$O_j = \prod_{n \in \mu_j} e^{-i\kappa_{2,n}t} \prod_{m \in \nu_j} e^{-\beta\kappa_{1,m}}, \tag{5}$$

which is invertible. With this definition, we have $h_j = h_j^\dagger \geq 0$ and $h_j|\Psi(\beta, t)\rangle = 0$.

We define the parent Hamiltonian of $|\Psi(\beta, t)\rangle$ as

$$H(\beta, t) = \sum_{j=1}^N h_j. \tag{6}$$

Note that we have now also suppressed the dependence of h_j on β and t for convenience.

Let us show that, indeed, (3) is the unique ground state of such an operator. Since $h_j|\Psi(\beta, t)\rangle = 0$, we have that $H(\beta, t)|\Psi(\beta, t)\rangle = 0$, and since $H(\beta, t) \geq 0$, this implies that $|\Psi(\beta, t)\rangle$ is a ground state of $H(\beta, t)$, with ground-state energy 0. Conversely, if $H(\beta, t)|\Psi'\rangle = 0$, then $h_j|\Psi'\rangle = 0$ and thus $\prod_{j \in \lambda(\kappa_{2,n})} e^{-i\kappa_{2,n}t} e^{-\beta\kappa_{1,m}} |\Psi'\rangle = 0$ for all j , which in turn means that $|\Psi'\rangle$ is proportional to $|\Psi(\beta, t)\rangle$: We thus see that $|\Psi(\beta, t)\rangle$ is the unique ground state of $H(\beta, t)$.

D. Adiabatic preparation

The existence of a smooth path of Hamiltonians connecting $H(\beta, t)$, and thus $|\Psi(\beta, t)\rangle$, to a simple product state at $H(0, 0)$ implies that these states can be prepared adiabatically, by starting with the product state $|\Psi(0, 0)\rangle = |\varphi_1\rangle \otimes \dots \otimes |\varphi_N\rangle$ and adiabatically changing the Hamiltonian from $H(0, 0)$ to $H(\beta, t)$. In fact, the first step of the procedure, changing the Hamiltonian from $H(0, 0)$ to $H(0, t)$, corresponds to applying a unitary transformation $U = e^{itK_2}$ to $|\Psi(0, 0)\rangle$. This transformation can be implemented *exactly* in time t by evolving with K_2 (rather than H). Alternatively, using the fact that K_2 is a sum of local commuting terms, U can be decomposed into a finite-depth local unitary circuit (where the number of layers only depends on the structure of the interaction), which can be realized exactly on a digital quantum computer or simulator in constant time.

The task that needs to be implemented adiabatically is the second part of the preparation, that is, the interpolation from $|\Psi(0, t)\rangle$ to $|\Psi(\beta, t)\rangle$. Generally, the time required for a faithful adiabatic evolution will depend on the magnitude of the spectral gap of $H(\beta', t)$ along the interpolation $\beta' \in [0; \beta]$. As it turns out, for the given type of interpolation, we can devise an efficient way of checking the presence of such a gap numerically, which we present in Sec. III A. Once we have established a lower bound on such a gap, we can use any standard bound for adiabatic evolutions [44,45], which gives an adiabatic runtime scaling as $T = O(N^2 \Delta^{-3} \epsilon^{-1})$, with Δ a lower bound on the gap, and ϵ the error in the final state.

Moreover, for the situation at hand, we can improve the scaling of the adiabatic preparation by making use of the locality of the Hamiltonian, combined with the version of the adiabatic theorem proven in Ref. [41]. To this end, we construct an alternative interpolation from $H(0, t)$ to $H(\beta, t)$ where, in each step, we only change *one* of the terms in the Hamiltonian; importantly, the method derived in Sec. III A to prove gaps still applies in that case. By changing the imaginary time in a suitably smooth way along this interpolation, we then obtain that the adiabatic time required per Hamiltonian term changed scales logarithmically with the desired accuracy. We can now concatenate the interpolation for all the

individual Hamiltonian terms. However, we would still have a Hamiltonian acting on the whole lattice, which will translate in an overhead in the runtime if we want to devise a digital algorithm. To overcome this, we use the fact that changes performed on distant terms can be equally well carried out in parallel due to an effective light cone through Lieb-Robinson bounds [41,46]. This means that at each step we only work with Hamiltonians supported on a region of size polylogarithmically in the system size. Combining all these results leads to a preparation scheme for $|\Psi(\beta, t)\rangle$ where the adiabatic time scales *polylogarithmically* with the desired accuracy and the problem size N , and thus exponentially better than other known methods for preparing MPS and PEPS. For a more technical discussion on this preparation method, we refer the reader to [41]. The bounds for the continuous-time version of the algorithm can essentially be derived from the ones presented in the original work.

E. Connection to tensor networks

Let us now show that the states (3) can be efficiently described as PEPS [17] with the same connectivity as $\mathcal{G}(\mathcal{V}, \mathcal{E})$ and a bond dimension D which is upper bounded by a function of r_1, r_2, z , and d , but which does not depend on N or M . To see this, let us consider an edge $(i, j) \in \mathcal{E}$. Since the operators $\kappa_{1,n}$ commute pairwise, we can express the product of operators that act on both i and j as a single one, $\prod_{i,j \in \lambda(\kappa_{1,n})} e^{it\kappa_{1,n}} = e^{it \sum_{i,j \in \lambda(\kappa_{1,n})} \kappa_{1,n}}$. We can bound the number of terms that appear in the sum in the exponent by the number of operators that act on qudit i . Since the individual operators $\kappa_{1,n}$ act on a radius r_1 , note that $e^{it \sum_{i \in \lambda(\kappa_{1,n})} \kappa_{1,n}}$ acts on, at most, all the qudits that are at a distance less or equal than $2r_1$ from qudit i . The number of such neighbors can be bounded by

$$x \leq z \sum_{i=1}^{2r_1-1} (z-1)^i = z \frac{1 - (z-1)^{2r_1}}{2-z} = O(z^{2r_1}). \tag{7}$$

Finally, note that an operator that acts on x qudits increase the bond dimension at most by $O(d^x)$.² Iterating this for every edge (which overcounts interactions) gives an upper bound $O(d^{z^{2r_1}})$ for the required bond dimension. Since a similar argument is valid for the operators $\kappa_{2,m}$, we conclude that the states of the form (3) can be described by a tensor network of bond dimension at most $O(d^{z^{2(r_1+r_2)}})$. Importantly, the bond dimension remains bounded independent of the system size since r_1, r_2 , and z are size independent.

In the following, we particularize the above results to standard PEPS. In particular, we will show that it contains all injective MPS and PEPS, as well as a PEPS corresponding to classical models, which (e.g., on the square lattice) are described by noninjective PEPS with a unique ground state. On the other hand, we generally cannot expect that G -injective PEPS, that is, those which can exhibit topological order, are of this form, regardless of boundary conditions and thus ground-space degeneracy, as they are not connected adiabatically

²This can be easily checked by iteratively performing a singular value decomposition on the operator and representing it as a matrix product operator (MPO) acting on the x qudits.




FIG. 2. MPS construction in terms of the operator $e^{i\kappa_{2,n}t}$, that creates an entangled pair between sites R_n and L_{n+1} , and $e^{\beta\kappa_{1,n}}$, that maps the virtual sites L_n, R_n to the physical qudit at position n .

to product states; however, this can be easily remedied by choosing a renormalization group (RG) fixed point in the corresponding phase as an initial state, rather than a product state, allowing for all findings in this paper to be carried over with only minor changes.

We will in the following consider regular lattices in one or higher dimensions. In the first case, we have MPS, and for the higher-dimensional case, we have PEPS.

1. Injective MPS

Matrix product states (MPS) are the simplest TN [17,18]. They can be written as

$$|\Psi\rangle = \sum_{s_1, \dots, s_N=0}^{d-1} A_{s_1}^1 \dots A_{s_N}^N |s_1, \dots, s_N\rangle. \tag{8}$$

We consider a special subclass, so-called injective MPS. They fulfill $d = d_0^2$, and are constructed with two qudits L_n and R_n each of dimension d_0 per vertex, and

$$|\Psi\rangle = \bigotimes_{n=1}^N Q_n |\Phi\rangle, \tag{9}$$

where $|\Phi\rangle$ is a state where, at each vertex, the qudit R_n is in a maximally entangled state with the qudit L_{n+1} on the vertex to its right (and thus L_n in a maximally entangled state with R_{n-1}), and $0 < Q_n \leq 1$ are invertible operators that transform $\mathbb{C}^{d_0} \otimes \mathbb{C}^{d_0} \rightarrow \mathbb{C}^d$. Generically, MPS become injective by blocking $\leq D^4$ original qudits [47].

MPS obviously fall within the family of states defined in (3) for the special case of a one-dimensional (1D) lattice as graph. Here, for each n , $\kappa_{2,n}$ acts on R_n and L_{n+1} in such a way that $e^{i\kappa_{2,n}t}$ creates the maximally entangled states on those qudits. The operator $\kappa_{1,n}$ is given by $\kappa_{1,n} = -\ln(Q_n)/\beta$ and acts on a single vertex each. This is illustrated in Fig. 2.

2. Injective PEPS

Injective PEPS are the generalization of (9) to higher spatial dimensions. The state has the same form, but now there are z_n qudits at site n , where z_n is the coordination number of vertex n (i.e., the number of edges connected to the vertex). The state $|\Phi\rangle$ contains entangled pairs along all possible vertices of the lattice. As for MPS, one can readily see that these states lie in the family defined in (3).

3. PEPS corresponding to classical models

Let us consider a classical model in a lattice. These ground states can be described by PEPS, which might be noninjective, but whose parent Hamiltonian has a unique ground state [48]. Consider a Gibbs state of some classical Hamiltonian $H_{cl}(s_1, \dots, s_n) = \sum_{\langle i_1 \dots i_k \rangle} h_{i_1 \dots i_k}(s_{i_1} \dots s_{i_k})$, where $s_i \in \{0, 1, \dots, d-1\}$, and $\langle i_1 \dots i_k \rangle$ denotes a the regions of neighboring particles coupled by the interaction. We define the state

$$|\Psi\rangle = \frac{1}{Z} \sum_{s_1, \dots, s_n} e^{-\beta H_{cl}(s_1, \dots, s_n)/2} |s_1, \dots, s_n\rangle, \quad (10)$$

where Z is the classical partition function. We can rewrite this state as follows [49]:

$$|\Psi\rangle = \frac{1}{Z} e^{-\beta \hat{H}_{cl}/2} \left(\frac{1}{\sqrt{d}} \sum_{s=0}^{d-1} |s\rangle \right)^{\otimes n}, \quad (11)$$

where \hat{H}_{cl} is an operator with eigenstates $|s_1, \dots, s_n\rangle$ and eigenvalues $H_{cl}(s_1, \dots, s_n)$. This state is of the type of Eq. (3), where $K_1 = H_{cl}/2$ and $K_2 = 0$. A description of these states in terms of PEPS can be easily obtained. For the special case of two-body nearest-neighbor interactions, we get a bond dimension equal to the dimension of the physical degree of freedom (see [49]).

III. GAPS AND EXPECTATION VALUES

In this section we establish our two key results, that find application in both efficient preparation protocols for the states (3) as well as for their verification. First, we develop an effective method to compute lower bounds to the gap of the parent Hamiltonians (6) for some range of parameters t, β . This will ensure that the corresponding ground states can be efficiently prepared with the adiabatic algorithm presented in the next section. Second, we determine the expectation values of a complete set of operators in those states, so that one can use them to check that the state has been successfully prepared.

A. Gaps

We will now explain how to efficiently obtain a lower bound on the gap of the parent Hamiltonians constructed in the previous subsection, for specific points in parameter space as well as uniform bounds for a whole parameter regime. To start with, note that $H(\beta = 0, t)$ is a sum of commuting projectors h_j for any t , and thus has a gap $\Delta = 1$. It thus remains to supply methods to bound the gap in the case where $\beta > 0$.

Let us first discuss how to obtain such a bound for a specific point $H(\beta, t) = \sum h_j$ in parameter space. To this end, consider the semidefinite program (SDP)

$$\delta = \max_{a_{ij}, c_{ij}} \quad (12a)$$

subject to

$$\forall i \neq j : h_i h_j + h_j h_i + a_{ij} h_i^2 + a_{ji} h_j^2 - c_{ij} h_i - c_{ji} h_j \geq 0, \quad (12b)$$

$$\forall i : \sum_{j \neq i} a_{ij} = 1, \quad (12c)$$

$$\forall i : \sum_{j \neq i} c_{ij} = x. \quad (12d)$$

For any feasible point of the SDP [and in particular the optimum in (12a)], summing (12b) over all $i \neq j$ yields (counting each pair twice, and up to a factor 2)

$$\sum_{i \neq j} h_i h_j + \sum_i h_i^2 - x \sum_i h_i \geq 0 \quad (13)$$

or, equivalently,

$$H^2 - xH \geq 0, \quad (14)$$

which says that H has no eigenvalues in the interval $(0; x)$. Since the ground-state energy is 0 by construction, this implies the existence of a spectral gap $\Delta \geq \delta$ above the unique ground state.

Since (12) is an SDP with the dimension of the constraints independent of N , it can be solved efficiently. Note that the SDP can be simplified considerably by setting

$$a_{ij} = c_{ij} = 0 \quad \text{for } \lambda(h_i) \cap \lambda(h_j) = \emptyset \quad (15)$$

and observing that (12b) is trivially satisfied in those cases, leaving a number of equations linear in N . Another possible simplification amounts to first solve for each *individual pair* $i \neq j$ the SDP which minimizes $a_{ij} + a_{ji}$ subject to $h_i h_j + h_j h_i + a_{ij} h_i^2 + a_{ji} h_j^2 \geq 0$; then, if for all i , there exists some $0 < y_i \leq 1$ such that $\sum_{j \neq i} a_{ij} < y_i$, then there is a gap [since h_i is relatively bounded by h_i^2 , so we can add positive contributions to both a_{ij} and c_{ij} while still satisfying (12b)]; a bound on the gap can either be computed directly from y_i or through the SDP (13) while keeping a_{ij} fixed. Finally, note that for $\beta = 0$ (where the h_i are commuting projectors), condition (12b) holds for any choice of $a_{ij} = c_{ij}$, and thus indeed gives $\delta = 1$.

Better bounds on the gap can be obtained by relaxing (12b) to only hold when summed over specific groups of terms [which still implies (13)]; a natural such case would be the variant of the SDP obtained by first grouping adjacent terms $\tilde{h}_i = \sum h_i$ (with the sum over terms in some neighborhood) and then setting up the SDP for the \tilde{h}_i . Furthermore, one can replace h_i (also after blocking) by projectors with the same range as h_i , since those relatively bound h_i and thus have a system-size independent gap if and only if the original Hamiltonian does.

Having described a way how to efficiently obtain a lower bound to the gap for a given point $H(\beta, t)$, how can we use this to build methods for certifying a gap over a whole range of parameters? The idea is based on the continuity of the SDP conditions (which should come as no surprise, given the finite dimension and smooth dependency of β of all objects involved). Given a certified gap δ for some $H(\beta, t)$ using the SDP (12) (with corresponding optimal parameters a_{ij}^* and c_{ij}^*),

we show in Appendix B that the SDP for $H(\beta + \tau, t)$ has a feasible point with $a'_{ij} = a^*_{ij}$, and where

$$c'_{ij} = e^{-2|v_j|\tau} c^*_{ij} - (1 - e^{-2|v_j|\tau}) - a^*_{ij}(e^{-2|v_j|\tau} - e^{-2|v_i\setminus v_j|\tau}) \tag{16}$$

for $0 \leq a_{ij} \leq 1$, and a variant thereof if $a_{ij} < 0$ or $a_{ij} > 1$ [see Eqs. (B13)]. Importantly, c'_{ij} changes uniformly continuous as τ is increased starting from $\tau = 0$. Thus, this allows one to obtain a lower bound $\delta(\tau)$ on the gap of $H(\beta + \tau, t)$ by virtue of (12d); importantly, the lower bound is uniform for a given interaction geometry and independent of the system size or the specific chosen model, and changes uniformly continuous with τ . We now take the τ_0 for which the lower bound closes, i.e., $\delta(\tau_0) = 0$, and rerun the SDP for $H(\beta + \tau_0, t)$: If that SDP returns a gap as well, this proves that $H(\beta + \tau, t)$ is gapped for all $0 \leq \tau \leq \tau_0$. By starting from $\beta = 0$ [for which (12) trivially holds], we can thus establish the existence of a gap (and obtain explicit lower bounds to it) for a finite regime $0 \leq \beta \leq \beta_0$ by evaluating the SDP (12) only at a finite number of points (where the bound can be improved by increasing the density of points). Note that for translational-invariant systems, we can show that the SDP from (12) can be “symmetrized” (see Appendix C). This means that for a given feasible point, with values of the variables (a^*_{ij}, c^*_{ij}) , we can find another feasible point by averaging a^*_{ij}, c^*_{ij} , such that all coefficients a_{ij} and c_{ij} are equal for all pairs of terms in (12a). Moreover, this solution provides the same value of δ . This directly shows that in this case δ does not change with the system size, thus recovering the result first showed in Ref. [18] about the existence of a gap independent of the system size. Therefore, for translational-invariant (TI) systems the size of the neighborhood τ is independent of the system size.

Let us now briefly discuss the suitability of the method for the TN of the previous section. For translational-invariant (or suitably uniform) injective MPS, it was proven by Fannes, Nachtergaele, and Werner [18] that the parent Hamiltonian is always gapped. Indeed, they showed that by blocking a number of sites proportional to the correlation length (for a fixed bond dimension), which in turn can be bounded as a function of the gap, and replacing the blocked Hamiltonian \tilde{h}_i by projectors, Eq. (12b) is fulfilled (with all a_{ij} and c_{ij} equal, as mentioned in the previous paragraph), and thus the SDP will yield a gap $\delta > 0$, already with the restriction (15) imposed. In higher dimensions, an analogous result has been obtained for PEPS which are unique ground states of local Hamiltonian (in particular, injective PEPS) [50]. Whenever the Hamiltonian is gapped in the thermodynamic limit, the SDP with condition (13) will be satisfied by the projector-valued Hamiltonian obtained after blocking a number of sites which only depends on the gap and the geometry of the system, but not on the system size or details of the model.

Finally, regarding the PEPS corresponding to classical models in dimensions higher than one, for sufficiently high temperatures (small β) the SDP method will give a gap (as the Hamiltonian at $\beta = \infty$ is trivial). Note that the corresponding classical model may have a phase transition at sufficiently low temperatures, which implies that the correlation length will diverge and thus the gap will vanish in the thermodynamic limit ($N \rightarrow \infty$). Thus, the SDP automatically also allows to

determine an upper bound to that critical temperature, which will yet again get more accurate as we consider larger regions, both by blocking and by relaxing (or omitting) the restriction (13). Thus, the continuity bound of Appendix B at the same time provides a means of determining upper bounds on the critical temperature of classical statistical models.

B. Expectation values

Computing expectation values of ground states of local Hamiltonians is hard in general [51]. However, for ground states of frustration-free Hamiltonians, there are certain observables for which it is straightforward to compute such values. In this section, we find a complete set of operators for which this can be done and which will be on the basis of the verification protocols presented below.

As argued in Sec. II C, for $|\Psi\rangle$ given in (3) we have $h_n|\Psi\rangle = 0$, and thus trivially $\langle\Psi|h_n|\Psi\rangle = 0$. We will now define a set of observables for which one can also compute the expectation value in $|\Psi\rangle$. We will restrict to qubits ($d = 2$) and will take $|\varphi_j\rangle = |0\rangle$, although it is straightforward to extend the method to qudits and other states. We will denote the Pauli operators acting on the j th qubit, $j = 1, \dots, N$, by σ_α^j with $\alpha = x, y, z$; for instance, $\sigma_z|0\rangle = |0\rangle$.

The key idea is to notice that for any $\lambda \subset \mathcal{V}$,

$$O_\lambda|\Psi\rangle = \left(\bigotimes_{j \in \lambda} |0\rangle_j \langle 0|_j\right) O_\lambda|\Psi\rangle, \tag{17}$$

where

$$O_\lambda = \prod_{n \in \mu(\lambda)} e^{-ik_{2,n}t} \prod_{m \in \nu(\lambda)} e^{-\beta\kappa_{1,m}} \tag{18}$$

with $\mu(\lambda) = \bigcup_{j \in \lambda} \mu_j$, $\nu(\lambda) = \bigcup_{j \in \lambda} \nu_j$, and where the sets μ_j, ν_j have been defined in Sec. II C.

The first set of operators is defined as $Z_\lambda^+ = (Z_\lambda + Z_\lambda^\dagger)/2$ and $Z_\lambda^- = (Z_\lambda - Z_\lambda^\dagger)/2i$, where

$$Z_\lambda = O_\lambda^{-1} \left(\bigotimes_{j \in \lambda} \sigma_z^j\right) O_\lambda, \tag{19}$$

that is, $|\Psi\rangle$ is a right (left) eigenvector of Z_λ (Z_λ^\dagger), using (17). We then have

$$\langle\Psi|Z_\lambda^+|\Psi\rangle = 1, \tag{20a}$$

$$\langle\Psi|Z_\lambda^-|\Psi\rangle = 0, \tag{20b}$$

and thus Z_λ^\pm have fixed expectation values.

The second set is more ample. Given any P supported in $\lambda \subset \mathcal{V}$ with the property that there exists $j \in \lambda$ such that $\langle 0|P|0\rangle_j = 0$, we define

$$Q_\lambda = O_\lambda^\dagger P O_\lambda. \tag{21}$$

Again, using (17) we can compute the expectation value of those operators

$$\langle\Psi|Q_\lambda|\Psi\rangle = 0. \tag{22}$$

The set of operators defined above, taken jointly for all $\lambda \in \mathcal{V}$, is complete in \mathcal{A} in the sense that their expectation values completely determine the state. To show that, we just have

to devise a subset of 4^N linearly independent operators from that set.

To this end, let us start from the set of all operators P which are a product of Pauli and identity operators, and associate to each of them an operator $Q \equiv Q(P)$ using one of the constructions above. Each of these operators can be written as

$$P = \bigotimes_{j \in \lambda} \sigma_{\alpha_j}^j, \quad (23)$$

where $\alpha_j = x, y, z$, and $\lambda \subseteq \mathcal{V}$ is the set of sites on which P acts nontrivially. In case $\alpha_j = z$ for all $j \in \lambda$, we define $Q(P) = Z_\lambda^+$; otherwise, $Q(P) = Q_\lambda$ [Eq. (21)]. Finally, for $P = \mathbb{1}$, let $Q(P) = \mathbb{1}$. In this way, starting from all products of Pauli operators and the identity, we have obtained a set of operators Q_n , $n = 1, \dots, 4^N$. This set is linearly independent iff the matrix $B_{n,m} = \text{tr}(Q_n^\dagger Q_m)/2^N$ is not singular. Trivially, for $\beta = 0$, $B_{n,m} = \delta_{n,m}$ and thus not singular. Since the operators O_λ used to define the map $P \mapsto Q$ are analytic in $\beta > 0$, the determinant of $B_{n,m}$ will be analytic as well, and thus it can only vanish at countably many points, all of which are isolated. Thus, generically it will be nonzero and, in the possible measure-zero cases where it is can be circumvented by taking a close-by value of β .

The fact that the set of operators Z_λ and Q_λ is (over)complete means that any observable can be expanded as a linear combination thereof. If we are able to obtain the corresponding coefficients, then we will be able to compute the expectation value of all physical quantities. In practice, this will be difficult since, due to the fact that the operators O_λ nontrivially overlap with each other, we will typically need an exponential number of terms in the expansion. Nonetheless, there may be a way of truncating that expansion, which would give lead to new algorithms in terms of tensor networks. Furthermore, since we know the expectation values of a basis of operators, we possess full tomographic information on state. However, as before, it is not useful to compute other expectation values. In Appendix D we show that the norms of the observables Q_λ decay at most exponentially with their support size, and thus they are guaranteed to have bounded (polynomially decaying) norm when the size of the support is fixed [at most $O(\ln N)$].

The operators Q_λ are supported on a set of vertices that are larger than λ (roughly speaking, on all vertices that lie at a distance up to $2r$ from that set). It is possible to define other observables which have a smaller support and for which we can still compute the expectation values. This is relevant for more practical applications, like the verification protocols introduced below, where we want to make statements about measurements performed within the support of Q_λ , and we want them to include as few qubits as possible. For that, given $j \in \mathcal{V}$ and an operator P' supported on λ with $j \notin \lambda$, we define

$$Q_{j,1} = O_j^\dagger \sigma_x^j P' O_j, \quad (24a)$$

$$Q_{j,2} = O_j^\dagger \sigma_y^j P' O_j, \quad (24b)$$

$$Q_{j,3} = O_j^\dagger (\mathbb{1} - \sigma_z^j) P' O_j, \quad (24c)$$

which only act on the joint support of O_j and P' . Again, using (17) we find the expectation value of those operators to be

$$\langle \Psi | Q_{j,\alpha} | \Psi \rangle = 0. \quad (25)$$

In particular, if we choose P' to be an arbitrary Pauli product, then $Q_{j,1}$ and $Q_{j,2}$ can be used to replace the operators Q_λ from Eq. (21) in our complete set of operators. At the same time, the operators $Q_{j,1}$ and $Q_{j,2}$ are still products of Pauli operators almost everywhere, except on the support of O_j which has a fixed size. This means that each $O_{j,1}$, $O_{j,2}$ can be estimated efficiently [i.e., with a number of (Pauli) measurements which only depend on the accuracy but not on the system size], and the only remaining operators in the complete set which cannot be estimated efficiently individually are the Z_λ^\pm .

In summary, we have defined a set of observables whose expectation values are either zero or one. We can choose a subset thereof where $|\lambda| \leq c$, with c a constant independent of N . In such a case, since we know O_λ we can efficiently write those observables as linear combinations of Pauli operators in the support of O_λ [like (19)], or even smaller [like (24)]. Note that the norms of these observables will also be efficient to compute in these cases, as we show in Appendix D.

IV. VERIFICATION SCHEMES

In this section we discuss different ways of exploiting the state preparation procedure for the state (3) as a verification scheme. First, we will analyze a quantum state verification method and show how to certify that the state has been created successfully by performing local measurements and using the fact that there exists a parent Hamiltonian that is both gapped and frustration free [52]. Then, we will consider the scenario of classical verification of quantum computation, where the goal is to make sure that someone else is in possession of a quantum computer solely through classical communication [21,22,29–32]. We will restrict here to the case of qubits, although the arguments can be easily extended to qudits.

A. Quantum state verification

Unique ground states of frustration-free Hamiltonians, like the ones we are dealing with here, can be trivially certified with local measurements. This is achieved by just performing local quantum tomographies to make sure that the expectation values $\langle h_j \rangle = 0$ for all h_n defined in (4). Indeed, if this is the case, then $\langle H \rangle = 0$ [cf. (6)], and since $H \geq 0$ and has a unique ground state, this implies that the measurement must have been performed on that state.

In practice, since we can only perform a finite number of measurements, the estimates for $\langle h_j \rangle$ will not be exactly zero; additionally, measurement errors will give rise to errors in those quantities as well. However, one can still estimate the success probability of the preparation in different ways. The most straightforward one is to relate the obtained expectation value of $\langle H \rangle = \text{tr}(H\rho)$ (with the corresponding estimated error) to the overlap of the state we have prepared ρ and the target ground state $|\Psi\rangle$. It is straightforward to show that the

fidelity obeys

$$\langle \Psi | \rho | \Psi \rangle \geq 1 - \frac{\text{tr}(H\rho)}{\Delta} \geq 1 - \frac{\text{tr}(H\rho)}{\delta}, \quad (26)$$

where Δ is the spectral gap of H , and δ the lower bound obtained in the previous section. Thus, if we can obtain this expectation value (with an error bar) and we compute δ , we directly obtain a lower bound to the overlap.

Neglecting measurement errors, for a finite number of measurements, the estimate of $\langle \Psi | \rho | \Psi \rangle$ will have some error bar. To use the bound (26), we need that the error in $\text{tr}(H\rho) = \langle H \rangle = \sum \langle h_j \rangle$ is sufficiently below δ . Assuming that we perform independent measurements and thus that we have independent errors with same standard deviation σ for the estimator of all the h_j , we have that the error ϵ_j for each term must satisfy $\delta/\sqrt{N} \sim \epsilon_j \sim \sigma/\sqrt{L_j}$, where L_j is the number of measurements performed to estimate $\langle h_j \rangle$. If, in addition, we assume $L_j = L_j$, we obtain a conservative estimate for the total number of measurements of $L_{\text{tot}} \lesssim NL_j \sim \sigma^2 N^2 / \delta^2$. Note that instead of performing independent measurements, one could measure qubits belonging to nonoverlapping regions in parallel [53], which might lead to a reduction of the total number of rounds.

The verification can also be analyzed as an adversarial game, where the prover prepares a state ρ , gives it to verifier, and claims that it is indeed Ψ . The verifier can then perform measurements to gain confidence that this is indeed the case. Such a protocol has been analyzed in Ref. [28] by assuming that the verifier can measure by projecting in the basis of eigenstates of the local operators of the Hamiltonian h_j . In Appendix E, we perform a similar analysis, but assuming that the verifier can only perform Pauli measurements, which might be a more realistic assumption for current experimental setups.

B. Classical verification of quantum computation

Let us now turn towards a different kind of verification, in which the verifier is fully classical and communicates with the quantum prover through a classical channel [21,22,29–32]. Such protocols can also be formulated as an adversary game: Here, the prover claims that he can efficiently carry out quantum computations on a quantum computer. The verifier has to make sure that this is the case by communicating classically with the prover, that is, asking him to perform certain tasks on his quantum computer and report the results. Of particular interest is the case where both prover and verifier have limited additional resources, namely, they can perform classical computations with a computational time that scales at most polynomially with the number of qubits. In this setting, the verifier can pose challenges to the prover which he can only accomplish if he has a quantum computer, but not with his limited classical resources, and the challenge is to find a way which allows the verifier with her limited classical resources to verify this.

Recently, several protocols achieving this task have been proposed whose security is based on standard complexity assumptions [29–31]. While the first protocol [29,30] is most adequate for fault-tolerant quantum computers, the latter [31] is very attractive since it can already be used to verify existing

noisy intermediate-scale quantum (NISQ) devices [54]. However, it requires the verifier to carry out classical computations whose runtime scales exponentially with the number of qubits, though with a relatively small exponent which makes it comparatively practical. Apart from their use to certify quantum computers that can be only used remotely by classical means, one of the most appealing applications of such protocols is in the context of certified random number generators [30].

Ground states of frustration-free quantum Hamiltonians that can be efficiently prepared, like the ones presented in this work, may offer an alternative way for this kind of verification; specifically, one can exploit the task of reproducing correctly the expectation values of the observables introduced in Sec. III B as a challenge for the prover. In the following, we will describe such verification protocols, and analyze their security and the underlying complexity theoretic assumptions in different regimes. As we will see, the straightforward application of this idea requires exponential time. More efficient versions of the protocol are possible, but the underlying complexity assumptions are less tangible and, thus, their security remains unclear.

1. Protocol

The verification protocol consists of three steps: (i) The verifier sends the prover instructions for preparing the state $|\Psi\rangle$. (ii) This step consists of R rounds: in each round, the verifier sends the prover a set of observables; the prover then prepares the state $|\Psi\rangle$, measures the observables, and reports the outcome. (iii) The verifier performs certain checks on the accumulated measurement outcomes to verify that the prover has indeed prepared the state $|\Psi\rangle$ and is thus in possession of a quantum computer.

For step (i), the verifier just has to give the circuit that prepares the state to the prover, which she can, e.g., obtain by trotterizing the adiabatic evolution. Alternatively, she can directly give the time-dependent Hamiltonian $H(t, \beta)$, together with the initial states $|\varphi_i\rangle$, to the prover, e.g., in case he is in possession of an analog quantum computer. An honest prover will then be able to efficiently prepare the state $|\Psi\rangle$ in a time that scales as $\ln(N)$, as discussed in Sec. II D.

For each round of step (ii), the verifier sends the prover a list of bases $\alpha = (\alpha_1, \dots, \alpha_N)$, $\alpha_j = x, y, z$, in which the individual qubits should be measured; the α will be generally drawn at random from some distribution which is dictated by the verification step (iii). The prover then prepares the state and measures qubit j in the Pauli basis α_j , and obtains results $s = (s_1, \dots, s_N)$, $s_j = \pm 1$. For an honest prover, these results are drawn from a distribution

$$P_0(s|\alpha) = \langle \Psi | \bigotimes_{j=1}^N \frac{1}{2} (\mathbb{1} + s_j \sigma_{\alpha_j}^j) | \Psi \rangle. \quad (27)$$

After receiving the measurement basis, the prover prepares $|\Psi\rangle$ and performs the measurements. Importantly, since $|\Psi\rangle$ can be prepared in time $O[\ln(N)]$, each round can in principle be carried out in time $\ln(N)$ as long as the prover has the ability to measure and communicate the results in parallel.

Step (ii) allows for several natural generalizations. In particular, we can allow for measurements beyond the Pauli basis, we can enable the verifier to make *adaptive* queries, where

the choice of \mathbf{s} in any round can depend on the results of the previous round, and we can split each round into several subrounds of communication, where the state is prepared once and then a sequence of adaptive measurements is performed on it, measuring only a subset of the qubits at each time.

In step (iii), the verifier uses the samples obtained from the prover to compute certain quantities which serve to verify that the prover is indeed sampling from the correct distribution P_0 . To this end, the verifier can, e.g., use some of the quantities Q constructed in Sec. III B for which $\langle Q \rangle = 0, 1$ (or variants thereof), or the Hamiltonian terms h_j for which $\langle \Psi | h_j | \Psi \rangle = 0$ (which can be used to replace the $Z_{(j)}^-$). Let us now consider one such observable Q supported on $\lambda \subset \mathcal{V}$. It can be decomposed as

$$Q = \sum_{\boldsymbol{\gamma}} o(\boldsymbol{\gamma}) \bigotimes_{j \in \lambda} \sigma_{\gamma_j}^j, \tag{28}$$

where $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_{|\lambda|})$, $\gamma_j = x, y, z$, and $o(\boldsymbol{\gamma})$ are the expansion coefficients. The verifier can then estimate $\langle Q \rangle$ using an estimator

$$\bar{Q} = \sum_{\boldsymbol{\gamma}} o(\boldsymbol{\gamma}) \bar{s}(\boldsymbol{\gamma}), \tag{29}$$

where $\bar{s}(\boldsymbol{\gamma})$ is the average outcome where the prover measured spin j in the basis σ_{γ_j} for all $j \in \lambda$, and an arbitrary basis for all $j \notin \lambda$; that is, if we denote the set of rounds where $\alpha_j = \gamma_j$ for all $j \in \lambda$ by $\mathcal{R}(\boldsymbol{\gamma})$, and the result of the r th round by $\mathbf{s}^r = (s_1^r, \dots, s_N^r)$, we have

$$\bar{s}(\boldsymbol{\gamma}) = \frac{1}{|\mathcal{R}(\boldsymbol{\gamma})|} \sum_{r \in \mathcal{R}(\boldsymbol{\gamma})} \prod_{j \in \lambda} s_j^r. \tag{30}$$

If the samples have been taken according to P_0 , $\bar{Q} \rightarrow \langle \Psi | O | \Psi \rangle$ in the limit $|\mathcal{R}(\boldsymbol{\gamma})| \rightarrow \infty$, which can be used as a means of verification (see Appendix E for a quantitative analysis). Note that, alternatively, we can determine the average also using only rounds $\mathcal{R}(\boldsymbol{\gamma})$ where sites $j \notin \lambda$ have only been measured in some specific bases.

As a simple example of this verification procedure, consider a resource state for measurement-based quantum computation (MBQC), such as the cluster state [55], or generalizations which allow for measurement-based computation using only Pauli measurements [56]. For those states, the preparation step is particularly easy, as they can be prepared by a single layer of commuting unitaries (i.e., $\beta = 0$) from a product state. They are frustration-free ground states of local Hamiltonians $H = \sum h_i \geq 0$, $h_i | \Psi \rangle = 0$, a property which allows for easy verification from the measurement statistics. And finally, we know that adaptive measurements in a suitable basis (Paulis and $\pi/4$ rotated states in the xy plane for the cluster state, or only Pauli measurements for the aforementioned generalization) are universal for quantum computation, making it impossible for a prover not in possession of a quantum computer to produce the correct distribution in such an interactive protocol, and at the same time imposing limitations on potential cheating strategies, as we will discuss in the following.

2. Complexity analysis

Let us now analyze under which conditions the protocol allows the verifier to conclude that the prover must indeed be in possession of a quantum computer, what potential classical cheating strategies might be, and what limitations to such strategies exist.

a. Sampling from the quantum distribution is hard. The first cheating strategy would be to find a way to classically sample from the correct distribution P_0 [Eq. (27)]; in that case, the verifier would have no way of detecting this. However, there are several strong complexity-theoretic arguments against that. First, note that, as already mentioned above, the adaptive version of the protocol contains measurement-based quantum computing: The cluster state is clearly in the given class (with $\beta = 0$), and an adaptive protocol with poly(N) queries per round would implement a general quantum computation. Thus, sampling from the correct P_0 is impossible unless $\text{BQP} = \text{BPP}$. However, it is also known that sampling from the output of a circuit of commuting gates with random product states as input and measuring in a fixed basis (a case which is contained in our protocol for $\beta = 0$) is computationally hard, as shown in [57]. In order to prove the hardness of this setup, in [57] they relate the complexity of sampling from such a circuit with the complexity of sampling from IQP circuits, a task that is known to be hard unless the polynomial hierarchy would collapse to its third level [58,59]. We thus conclude that there is complexity-theoretic evidence that it is impossible for the prover to classically sample from the correct distribution P_0 (up to constant error) in polynomial time.

b. Reproducing all $\langle Q \rangle$ is hard. The second cheating strategy would be to sample from some different distribution P' which is chosen such that all estimators for $\langle Q \rangle$ computed from P' are correct. Any estimator \bar{Q} [Eq. (29)] supported on a subset λ of sites can be computed in many different ways, which differ by the measurement settings over which we average for the sites not contained in λ [see Eq. (30) and the discussion below it]. We will thus additionally impose that the estimator \bar{Q} converges to the same value for all those ways to compute it; this can be easily ensured by the verifier by computing \bar{Q} in all different ways, or just in a randomly chosen one. (Alternatively, this can be ensured by computing the marginal probability distributions in different ways and checking their consistency, that is, by checking that the measurement setting α_i of qubit i does not affect the distribution of the measurement results s_j for any of the other qubits $j \neq i$; this can be seen as a kind of nonsignaling condition on the distribution.)

These conditions, however, imply that $P' = P_0$, the correct distribution derived from the quantum state. The reason is that for any given $\boldsymbol{\alpha}$, we can reconstruct $P'(\cdot | \boldsymbol{\alpha})$ from the expectation values of all Pauli measurements given by arbitrary substrings of $\boldsymbol{\alpha}$ [that is, where the Paulis in some positions have been replaced by identities, as in Eq. (23)] through an N -fold Hadamard transform (using the consistency condition above); the latter, in turn, can be reconstructed from all $\langle Q \rangle$ as they are related by an invertible transformation, as discussed in Sec. III B.

We thus find that this is not a viable cheating strategy since no other distribution which yields all the correct $\langle Q \rangle$ exists. Note that together with the hardness results mentioned under Sec. IV B 2 a, this implies that even sampling from a P' which approximates $\langle Q \rangle$ for all Q sufficiently well is hard. The required accuracy in Q has to be chosen such that the expectation values of Pauli strings have exponential accuracy in N ; since the number of samples required to this end is in fact determined by the latter [the transformation (29) is exact], this generally requires an exponential number of samples.

Let us now discuss two different cases for the verification protocol and their security. First, we consider the case where both prover and verifier have bounded (polynomial) storage space, but the verification procedure can take an exponential number of rounds, and the verifier can use exponential time. In this case, the arguments from before regarding the hardness of sampling from P_0 or reproducing the $\langle Q \rangle$ imply the security of the protocol. Second, we consider the case in which the verifier only tests local observables, i.e., those supported on a region of at most $O(\ln N)$. Note that the verification is efficient in this case, i.e., it can be carried out in polynomial time for a prover and verifier having both polynomial storage space as before. In this case, we also argue how successful cheating strategies will be likely to fail. Most interesting, if such a cheating strategy succeeded, it would imply the existence of classical algorithms for computing local observables for generic tensor network states.

c. A secure space-bounded and exponential-time verification protocol. In each round, the prover is only granted $\text{poly}(N)$ time (or logarithmic time with suitable parallel processing power). By performing an exponential number of queries, the verifier can get an exponentially precise estimate \tilde{Q} for any of the observables Q (either a randomly selected one, or all of them); importantly, the expansion coefficient $o(\gamma)$ of Q in the Pauli basis [Eq. (28)] is a trace of a product of local operators and can thus be computed in PSPACE, and the summands in (29) can be sampled and added sequentially one by one; the verifier thus only requires polynomial storage space. The prover, on the other hand, cannot classically sample from the correct distribution in the available polynomial (or logarithmic) time, due to the hardness results discussed in Sec. IV B 2 a; at the same time, his limited memory prevents him from precomputing all possible outcomes after having learned the adiabatic circuit for preparing $|\Psi\rangle$ in step (i) (even though the exponential time used for the overall protocol would allow for it).

Let us note that what makes our setup special is not the fact that the knowledge of all $\langle Q \rangle$ allows to reconstruct the probability distribution; such a complete tomography is possible in any scenario. Rather, what makes our construction special is that the verifier knows the required expectation values for all operators $\langle Q \rangle$ right away, whereas in a general tomography scheme, a computationally costly reconstruction procedure is required to obtain P from measured expectation values. However, in order to compute a general expectation value $\langle Q \rangle$, the verifier still needs to collect an exponential amount of data from the prover: While it is sufficient to sample a small number of randomly chosen Q 's, most Q 's still have an exponential number of terms in their Pauli expansion

(28). The situation here can thus be somehow regarded as the reverse from that introduced in Ref. [31]: While in that case the verifier requires a polynomial number of measurements from the prover, it takes her an exponential time to check if they correspond to the correct probability (by computing the cross entropy). In our case, the correct expectation values can be trivially computed, but one requires an exponential number of samples to obtain them. Note that computing the coefficients $o(\gamma)$ of Q , which are needed to estimate its expectation value from the samples, might require exponential resources as well.

d. Proposal for an efficient verification protocol and its security implications. As we have seen, checking some $\langle Q \rangle$, chosen at random from the set of all Q , allows to verify that the prover is in possession of a quantum computer. This, however, requires exponential resources, since a typical Q involves a significant fraction of the Pauli terms acting on its support, that is, an exponential number, and moreover we require an exponential accuracy. A practical verification protocol must thus resort to testing only *local* Q 's, that is, those which are supported on at most order $\ln(N)$ sites, to an accuracy $1/\text{poly}(N)$, as those can be sampled in $\text{poly}(N)$ time.

Thus, the question which arises is whether there is a way for the prover to efficiently sample from a distribution P' which reproduces those local $\langle Q \rangle$ correctly, up to polynomial accuracy.

First, let us notice here that, if the prover indeed measures in the Pauli basis, reaching a polynomial accuracy on local $\langle Q \rangle$ expectation values uniquely identifies the ground state $|\Psi\rangle$ among all multiqubit states. The reason is that we can include the terms of the parent Hamiltonian $H = \sum h_j$ in the set of Q 's, and since $|\Psi\rangle$ is the unique ground state of H and H is gapped, this implies $\| |\Phi\rangle - |\Psi\rangle \| < 1/\text{poly}(N)$ if the $\langle h_j \rangle$ are $1/\text{poly}(N)$ -close [see Eq. (26)]. Hence, to entirely establish the security of the protocol, it would be enough to prove that any distribution P' reproducing $\langle Q \rangle$ with the desired accuracy must be obtainable by performing Pauli measurements on a multiqubit state. Notice that a possible way to achieve that would be to prove a robust self-testing statement based on the $\langle Q \rangle$ expectation values (see [60] for a review of self-testing techniques).

Second, even if such a distribution P' exists, there are constraints on the design of algorithms to sample from it efficiently. For instance, one might assume that it should be possible to characterize the space of solutions of the equations $|\tilde{Q} - \langle Q \rangle| \leq 1/\text{poly}(N)$, which locally constrain P' , use them to find ways to sample locally correctly, and patch those ways of sampling together to obtain a sampling algorithm which works globally. However, such a strategy, if based only on the final conditions on $\langle Q \rangle$, is bound to fail, unless further properties of P_0 are taken into account. Specifically, we can choose the Q to enforce 3-SAT clauses or some other classical NP-hard problem [such as spin glasses on a two-dimensional (2D) lattice], in which case such an algorithm is bound to fail as it would have to solve the NP problem. (It is a fingerprint of hard instances of NP problems that the local constraints cannot be patched together easily.) Note that also knowing the precise local reduced density matrices does not necessarily make the problem easier, as the general quantum marginal problem is QMA-hard [61,62].

Thus, a successful cheating strategy most likely will have to use the full knowledge of the adiabatic path used to prepare the state or, equivalently, knowledge of K_1 and K_2 . Note that Osborne showed that the computation of local expectation values on states that undergo an adiabatic evolution under a local gapped Hamiltonian can be achieved classically [63]. However, this method does not help the prover either since it does not provide samples from the full probability distribution, which is what the prover is asked to return. Another approach could be to classically simulate the adiabatic evolution: for instance, one could try to adapt the Monte Carlo algorithm by Bravyi and Terhal [64] for the simulation of adiabatic quantum computation along a path which is both frustration free (which we have) and sign-problem free (which we do not necessarily have), followed by a measurement in the σ_z basis (which we do not have). Indeed, it is plausible that such an algorithm will allow to correctly sample in cases where the final Hamiltonian is classical and only needs to be sampled in the σ_z basis. On the other hand, it will likely break down if one of the conditions is not met, which will manifest itself in a sign problem in the Monte Carlo method. In fact, we cannot expect a cheating strategy based on simulating the adiabatic evolution to work since such an algorithm (if working as desired) will precisely sample from the *correct* distribution P_0 , which we have previously established to be computationally hard. Thus, in order to attack the protocol with such a strategy, one would have to devise a method to simulate the adiabatic evolution in a way where local expectation values are reproduced correctly, yet global properties would not (with the goal of circumventing hardness results); it is unclear how a route to accomplish this would look like.

ACKNOWLEDGMENTS

We thank A. Bouland, Z. Landau, and U. Vazirani for helpful discussions. This work has received support from the European Union’s Horizon 2020 program through the ERC-AdG QUENOCOBA (Grant No. 742102), the ERC-CoG SEQUAM (Grant No. 863476) and SuperQuLAN (Grant No. 899354), and the European Union’s Horizon Europe program through the ERC StG FINE-TEA-SQUAD (Grant No. 101040729), from the DFG (German Research Foundation) under Germany’s Excellence Strategy (Grant No. EXC2111 390814868), from the Austrian Science Fund (FWF) through Project No. 414325145 within SFBF7104, from the German BMBF through EQUAHUMO (Grant No. 13N16066) within the funding program “Quantum technologies - from basic research to market” and from the Bavarian Government through the Munich Quantum Valley. J.T. thanks the Alexander von Humboldt foundation and the Google Research Scholar Program for support.

APPENDIX A: EXAMPLES OF SET OF OPERATORS $\mathbb{K}_{r,M}$

In this Appendix we give some examples on how to construct operators $K \in \mathbb{K}_{r,M}$ defined in (2). Recall that K is defined as a sum of local operators $\kappa_n \in \mathcal{A}_r$, for $n = 1, \dots, M$, such that the operators κ_n commute pairwise.

Let us take a set of pairwise commuting unitary operators acting on at most r_U sites: $\{U_\alpha : |\lambda(U_\alpha)| \leq$

$r_U, [U_\alpha, U_\beta] = 0\}$. Note that a straightforward way of constructing commuting, Hermitian operators is as follows:

$$\kappa_n = \left(\prod_{\alpha:n \in \lambda(U_\alpha)} U_{\alpha_i} \right) O_n \left(\prod_{\alpha:n \in \lambda(U_\alpha)} U_{\alpha_i} \right)^\dagger, \quad (A1)$$

where O_n is a one-body operator acting on site n such that $\|O_n\|_\infty \leq 1$ [note that we can also take O_n to act on disjoint subset of vertices $\lambda(O_n) \subset \mathcal{V}$, such that $\lambda(O_n) \cap \lambda(O'_n) = \emptyset$]. Note that the product runs over all unitary operators whose support intersect site n . In this case, the final operators κ_n are labeled by a physical site, though in the general definition (2) this is not necessarily the case. Note that the conjugation by unitaries preserves the commutativity, i.e., $[\kappa_n, \kappa_m] = 0$ for all $n \neq m$, as well as the spectrum of the operators, so $\|\kappa_n\|_\infty \leq 1$. We can also upper bound the support size of κ_n , $|\lambda(\kappa_n)| \leq r_U - 1$.

It follows that it suffices to find a set of pairwise commuting unitary operators in order to construct an operator $K \in \mathbb{K}_{r,M}$. We will explicitly outline three different methods on how to construct such a set.

Diagonal unitaries. The first obvious way to do so is to take all the unitary operators diagonal on the same basis. Note that one can use these unitary operators to construct graph states [65,66] (such as the cluster state [65]) and weighted graph states [67] by applying them over a lattice where each edge represents a maximally entangled state between adjacent vertices. Thus, the family of states (3) trivially contains those.

Toric-code-type unitaries. An alternative approach for two-dimensional rectangular lattices is as follows. First, take some Hermitian operators h_A and h_B acting on adjacent plaquettes A and B and assume that $h_{A,B}$ are of a toric-code type [68] [see Fig. 3(a)]. By this we mean that $h_A = \bigotimes_{i \in A} O_A$, $h_B = \bigotimes_{i \in B} O_B$, where $O_{A,B} = O_{A,B}^\dagger$ and such that $\{O_A, O_B\} = 0$ (for the toric code we have $O_A = \sigma_z$, $O_B = \sigma_x$). Since the supports of h_A and h_B intersect in two vertices, the anticommutativity condition implies that $[h_A, h_B] = 0$. Thus, we can generate commuting unitary operators by evolving $h_{A,B}$ up to different times:

$$U_\alpha^A = e^{ih_A t_\alpha}, \quad U_\beta^B = e^{ih_B t_\beta}. \quad (A2)$$

By construction, we have $[U_\alpha^A, U_{\alpha'}^A] = [U_\alpha^A, U_\beta^B] = [U_\beta^B, U_{\beta'}^B] = 0$ for all $\alpha, \alpha', \beta, \beta'$. Note that we can choose $O(N)$ free parameters, corresponding to the different evolution times $\{t_\alpha\}_\alpha$ and $\{t_\beta\}_\beta$, as many as plaquettes in the lattice.

Bravyi-type unitaries. Finally, we present a more general approach, which is a generalization of the results of Bravyi *et al.* [69] about the characterization of commuting, local Hamiltonians in 1D (see also [70] for a discussion on this topic). We will present the method for the particular case of a 2D rectangular locality structure, but the generalization to other geometries is straightforward.

Let us start with a state defined over a lattice. At each site, we have four virtual particles of dimension D , which are mapped into a single physical qudit, with a physical degree of freedom d . For a visualization see Fig. 3(b), where we associate black dots with virtual particles; the big blue circles represent the physical sites, and correspond to the vertices of

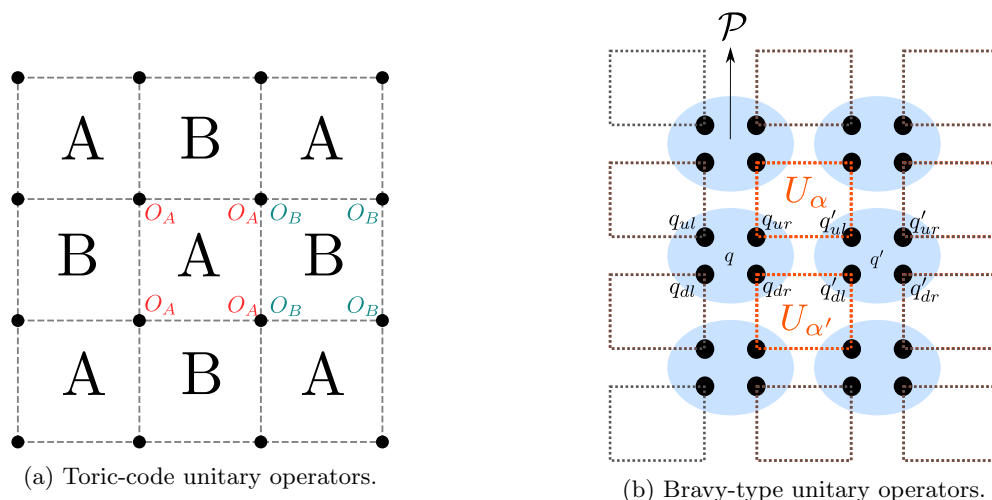


FIG. 3. (a) Representation of a toric-code Hamiltonian, in which two types of plaquette operators alternate. (b) Representation of a state that is obtained by mapping four virtual particles (black dots) into a physical one (blue shape). The shaded lines represent unitary operators acting on adjacent virtual sites, each of them belonging to a different qudit.

the lattice. We denote by $\mathcal{P} : (\mathbb{C}^d)^{\otimes 4} \mapsto \mathbb{C}^D$ the map from the virtual to the physical degrees of freedom. We take \mathcal{P} to be unitary, which means that $d^4 = D$.

Let us denote by $\mathcal{H}_q \cong \mathbb{C}^d$ the Hilbert space associated with the physical particle at vertex q . Bravyi and Vyalii [69] showed that \mathcal{H}_q can be decomposed as

$$\mathcal{H}_q = \bigoplus_i (\mathcal{H}_{q_{ul}}^i \otimes \mathcal{H}_{q_{ur}}^i \otimes \mathcal{H}_{q_{dl}}^i \otimes \mathcal{H}_{q_{dr}}^i), \quad (A3)$$

where $q_{ul}, q_{ur}, q_{dl}, q_{dr}$ correspond to the virtual degrees of freedom [see Fig. 3(b)]. We denote by $\mathcal{H}_{q_{ul}}$ the Hilbert space associated with the virtual particle q_{ul} and by $\mathcal{H}_{q_{ul}}^i$ a subspace of $\mathcal{H}_{q_{ul}}$.

Let us now define some unitary operators U_α , acting on four adjacent virtual particles, each belonging to a different vertex as in Fig. 3(b). Here α just denotes an index to enumerate the different operators, without any spatial meaning. We do not impose any restriction on these operators apart from the virtual particles in which they act. Let us now show that when we consider the operators U_α acting on the joint Hilbert space of two neighboring physical sites, i.e., on $\mathcal{H}_q \otimes \mathcal{H}_{q'}$, they commute. To see this, consider the examples from Fig. 3(b). We can write the operators U_α and U'_α as

$$\begin{aligned} U_\alpha | \mathcal{H}_q \otimes \mathcal{H}_{q'} &= o_{q_{ur}q'_{ul}} \otimes \mathbb{1}_{q_{ul}q_{dr}q_{dl}} \otimes \mathbb{1}_{q'_{ur}q'_{dl}q'_{dr}}, \\ U'_\alpha | \mathcal{H}_q \otimes \mathcal{H}_{q'} &= o_{q_{dr}q'_{dl}} \otimes \mathbb{1}_{q_{ul}q_{ur}q_{dl}} \otimes \mathbb{1}_{q'_{ur}q'_{ul}q'_{dr}} \end{aligned} \quad (A4)$$

for some operators $o_{q_{ur}q'_{ul}}$ and $o_{q_{dr}q'_{dl}}$ that act nontrivially only on the subspaces $\mathcal{H}_{q_{ur}} \otimes \mathcal{H}_{q'_{ul}}$ and $\mathcal{H}_{q_{dr}} \otimes \mathcal{H}_{q'_{dl}}$, respectively. With this decomposition, it is easy to see that $[U_\alpha, U'_\alpha] = 0$.

Now note that the application of the unitary operators \mathcal{P} does not affect the commutation; it simply changes the basis in which the qudits are expressed. Hence, we can define the final set of unitaries to be

$$\left(\prod_{\mathcal{H}_q \cap \lambda(U_\alpha) \neq \emptyset} P_q^\dagger \right) U_\alpha \left(\prod_{\mathcal{H}_q \cap \lambda(U_\alpha) \neq \emptyset} P_q \right), \quad (A5)$$

where $\lambda(U_\alpha)$ denotes as before the support of the operator U_α . Note that this defines a set of unitary operators over the full lattice, such that all of the operators commute pairwise.

Note that this technique does not cover the generation of states where $d = 2$. Therefore, for qubits other approaches like the ones presented above must be employed.

APPENDIX B: CONTINUITY BOUND ON THE GAP

Here, we show that the validity of condition (12b) for h_i [Eq. (4)] at a given point $H(\tau, t)$ (with τ the imaginary time formerly known as β) implies the validity of the same condition at a different point (τ', t) , with $\Delta\tau = \tau' - \tau > 0$, with Hamiltonian terms h'_i , only with different values c'_{ij} and c'_{ji} . Recall (12b):

$$\forall i \neq j : 0 \leq h_i h_j + h_j h_i + a_{ij} h_i^2 + a_{ji} h_j^2 - c_{ij} h_i - c_{ji} h_j. \quad (B1)$$

Let us pick a specific pair $i \neq j$. Let

$$\begin{aligned} A &= \prod_{m \in v_i \setminus v_j} e^{-\Delta\tau \kappa_{1,m}}, & C &= \prod_{m \in v_j \setminus v_i} e^{-\Delta\tau \kappa_{1,m}}, \\ B &= \prod_{m \in v_i \cap v_j} e^{-\Delta\tau \kappa_{1,m}} \end{aligned} \quad (B2)$$

be the part of the deformation $e^{(\Delta\tau)K_1}$ acting only on h_i , only h_j , and both h_i and h_j , respectively, that is,

$$h'_i = ABh_iAB, \quad h'_j = BCh_jBC. \quad (B3)$$

By construction,

$$\alpha \mathbb{1} \leq A^2 \leq \mathbb{1}, \quad \beta \mathbb{1} \leq B^2 \leq \mathbb{1}, \quad \gamma \mathbb{1} \leq C^2 \leq \mathbb{1}, \quad (B4)$$

where (unless we have tighter bounds)

$$\alpha = e^{-2|v_i \setminus v_j| \Delta\tau}, \quad \gamma = e^{-2|v_j \setminus v_i| \Delta\tau}, \quad \beta = e^{-2|v_i \cap v_j| \Delta\tau} \quad (B5)$$

(in particular, $\alpha, \beta, \gamma > 0$). Henceforth, for convenience of notation we will write $P := h_i$ and $Q := h_j$, as well as $a := a_{ij}, b := a_{ji}, c := c_{ij}, d = c_{ji}$.

Let us now derive two inequalities. First,

$$\leq (P + Q)B^2(P + Q) - \beta P^2 - \beta Q^2 \quad (B7)$$

$$= PB^2Q + QB^2P + PB^2P - \beta P^2$$

$$\beta(PQ + QP) = \beta(P + Q)^2 - \beta P^2 - \beta Q^2 \quad (B6)$$

$$+ QB^2Q - \beta Q^2. \quad (B8)$$

Second, using that $\mathbb{1} \leq C^{-2}$, $-\mathbb{1} \leq -\gamma C^{-2}$, $-A^2 \leq \alpha \mathbb{1}$, and $B^2 \leq \mathbb{1}$, we have that, for now assuming $0 \leq a \leq 1$,

$$PB^2P - \beta(1 - a)P^2 - \beta cP = aPA^2B^2P \otimes C^{-2} - aPA^2B^2P \otimes C^{-2} + PB^2P - \beta(1 - a)P^2 - \beta cP \quad (B9)$$

$$\leq aPA^2B^2P \otimes C^{-2} + [(1 - \alpha\alpha)PB^2P - \gamma\beta(1 - a)P^2 - \gamma\beta cP] \otimes C^{-2} \quad (B10)$$

$$\leq aPA^2B^2P \otimes C^{-2} + \underbrace{\{[(1 - \alpha\alpha) - \gamma\beta(1 - a)]P^2 - \gamma\beta cP\}}_{=:X} \otimes C^{-2} \quad (B11)$$

$$\leq aPA^2B^2P \otimes C^{-2} - c'P \otimes C^{-2}, \quad (B12)$$

where we have used that the term labeled X in (B11) is non-negative [as $(1 - \alpha\alpha) > (1 - a)$ and $\gamma\beta < 1$] to bound $XP^2 \leq XP$, and defined

$$c' := \gamma\beta c - (1 - \gamma\beta) - a(\gamma\beta - \alpha) \quad \text{for } 0 \leq a \leq 1. \quad (B13a)$$

The same inequality can be shown to hold in the other cases, with

$$c' := \gamma\beta c - (1 - \gamma\beta) - a(\gamma\beta - 1) \quad \text{for } a < 0; \quad (B13b)$$

$$c' := \gamma\beta c - (1 - \beta) - a(\beta - \alpha) \quad \text{for } a > 1, 1 - \alpha\alpha \geq 0; \quad (B13c)$$

$$c' := \gamma\beta c - (1 - \beta) - a\beta(1 - \alpha) \quad \text{for } 1 - \alpha\alpha < 0. \quad (B13d)$$

The same way, we can also obtain for $0 \leq b \leq 1$

$$QB^2Q - \beta(1 - b)Q^2 - \beta dQ \leq bA^{-2} \otimes QB^2C^2Q - d'A^{-2} \otimes Q, \quad (B14)$$

where

$$d' := \alpha\beta d - (1 - \alpha\beta) - b(\alpha\beta - \gamma) \quad (B15)$$

(or correspondingly, with α and γ exchanged with respect to c' , when $b < 0$ or $b > 1$). Starting from (B1), we thus obtain

$$0 \stackrel{(B1)}{\leq} \beta(PQ + QP + aP^2 + bQ^2 - cP - dQ) \quad (B16)$$

$$\stackrel{(B8)}{\leq} (PB^2Q + QB^2P + PB^2P - \beta P^2 + QB^2Q - \beta Q^2) + \beta(aP^2 + bQ^2 - cP - dQ) \quad (B17)$$

$$\stackrel{(B12, B14)}{\leq} PB^2Q + QB^2P + aPA^2B^2P \otimes C^{-2}bA^{-2} \otimes QB^2C^2Q - c'P \otimes C^{-2} - d'A^{-2} \otimes Q. \quad (B18)$$

Conjugating both sides of the equation with ABC , we now immediately see that this is nothing but the condition (B1) for the Hamiltonian h'_k , with a_{ij} and a_{ji} unaltered, and new feasible points c'_{ij} and c'_{ji} as given by Eqs. (B13) and (B15).

Importantly, in case $0 \leq a_{ij} \leq 1$ for all i, j , using (B5) this yields a feasible point with

$$c'_{ij} = e^{-2|v_j|\Delta\tau} c_{ij} - (1 - e^{-2|v_j|\Delta\tau}) - a_{ij}(e^{-2|v_j|\Delta\tau} - e^{-2|v_i\setminus v_j|\Delta\tau}), \quad (B19)$$

for all c_{ij} in (12b), which immediately implies a continuity bound on the gap which only depends on the geometry of the model, but not on the lattice size or the specifics of the model at hand. This can be straightforwardly adapted to the case where some a_{ij} are negative or above 1.

As outlined in the main text, this bound can be subsequently used, starting from $\beta = 0$, to determine a regime for which the gap can be lower bounded by a uniform $\Delta_0 > 0$. Clearly, tighter bounds can be obtained by choosing more intermediate values of β .

APPENDIX C: SYMMETRIZATION OF THE SDP FOR TRANSLATIONAL-INVARIANT SYSTEMS

Let us see that for translational-invariant systems, the SDP presented in the main text gives a bound δ which is independent of the system size. For simplicity, we consider 1D systems and the case in which Eq. (15) holds. The most general case, as well as the higher-dimensional case, can be generalized from this proof.

Let N denote the system size and let us consider periodic boundary conditions. Consider the SDP (12). Let us assume that we have a feasible point for some choice of the variables (a_{ij}^*, c_{ij}^*) . Let us show that we can redefine these variables, such as we still have a feasible point with the same value of δ , as follows:

$$\tilde{a}_{i,i+1} := \frac{1}{N} \sum_i a_{i,i+1}^*; \quad \tilde{a}_{i,i-1} := \frac{1}{N} \sum_i a_{i,i-1}^*, \quad (C1)$$

$$\tilde{c}_{i,i+1} := \frac{1}{N} \sum_i c_{i,i+1}^*; \quad \tilde{c}_{i,i-1} := \frac{1}{N} \sum_i c_{i,i-1}^*. \tag{C2}$$

It is easy to see that the constraints (12b) are still satisfied. Indeed,

$$\begin{aligned} & h_i h_{i+1} + h_{i+1} h_i + \tilde{a}_{i,i+1} h_i^2 + \tilde{a}_{i+1,i} h_{i+1}^2 - \tilde{c}_1 h_i - \tilde{c}_2 h_{i+1} \\ &= \frac{1}{N} \left[N(h_i h_{i+1} + h_{i+1} h_i) + \left(\sum_{k=1}^N a_{k,k+1}^* \right) h_i^2 + \left(\sum_{k=1}^N a_{k+1,k}^* \right) h_{i+1}^2 - \left(\sum_{k=1}^N c_{k,k+1}^* \right) h_i - \left(\sum_{k=1}^N c_{k+1,k}^* \right) h_{i+1} \right]. \end{aligned}$$

Now note that, because of the translational invariance, we have that the terms $h_i h_{i+1}$, $h_{i+1} h_i$, $h_i = h_i \otimes \mathbb{1}$ and $\mathbb{1} \otimes h_{i+1}$ are the same for all i . Thus, we can write $N(h_i h_{i+1} + h_{i+1} h_i) = \sum_{k=1}^N h_k h_{k+1} + h_{k+1} h_k$, and get

$$\begin{aligned} & h_i h_{i+1} + h_{i+1} h_i + \tilde{a}_{i,i+1} h_i^2 + \tilde{a}_{i+1,i} h_{i+1}^2 - \tilde{c}_{i,i+1} h_i - \tilde{c}_{i+1,i} h_{i+1} \\ &= \frac{1}{N} \sum_{k=1}^N \left[\underbrace{h_k h_{k+1} + h_{k+1} h_k + a_{k,k+1}^* h_k^2 + a_{k+1,k}^* h_{k+1}^2 - c_{k,k+1}^* h_k - c_{k+1,k}^* h_{k+1}}_{\geq 0} \right] \geq 0. \end{aligned}$$

For the constraints (12c) and (12d) we get

$$\begin{aligned} \sum_{j \neq i} \tilde{a}_{ij} &= \tilde{a}_{i,i+1} + \tilde{a}_{i,i-1} = \frac{1}{N} \sum_{i=1}^N a_{i,i+1}^* + \frac{1}{N} \sum_{i=1}^N a_{i,i-1}^* = \frac{1}{N} \sum_{i=1}^N \left(\underbrace{a_{i,i+1}^* + a_{i,i-1}^*}_{=1} \right) = 1, \\ \sum_{j \neq i} \tilde{c}_{ij} &= \tilde{c}_{i,i+1} + \tilde{c}_{i,i-1} = \frac{1}{N} \sum_{i=1}^N c_{i,i+1}^* + \frac{1}{N} \sum_{i=1}^N c_{i,i-1}^* = \frac{1}{N} \sum_{i=1}^N \left(\underbrace{c_{i,i+1}^* + c_{i,i-1}^*}_{=x} \right) = x. \end{aligned}$$

With the new variables \tilde{a}_{ij} , \tilde{c}_{ij} we actually have redundant information in (12b) and it suffices to write a single condition:

$$h_i h_{i+1} + h_{i+1} h_i + \tilde{a}_{i,i+1} h_i^2 + \tilde{a}_{i+1,i} h_{i+1}^2 - \tilde{c}_{i,i+1} h_i - \tilde{c}_{i+1,i} h_{i+1} \geq 0, \tag{C3}$$

$$\tilde{a}_{i,i+1} + \tilde{a}_{i,i-1} = 1, \tag{C4}$$

$$\tilde{c}_{i,i+1} + \tilde{c}_{i,i-1} = x. \tag{C5}$$

It is easy to see now that adding a new particle, i.e., $N \rightarrow N + 1$, does not change the SDP [since it will only add redundant information to (C3)]. Thus, given that for a TI finite system our method provides a nontrivial bound for the gap, the same bound holds for $N \rightarrow \infty$. In this case, we recover the result of the martingale method by Ref. [18]. Note that for higher dimensions, instead of a single condition (C3) we would have as many as different types of overlaps within the different Hamiltonian terms.

APPENDIX D: BOUND ON THE NORMS OF THE OBSERVABLES Q_λ

Let us consider the operators Q_λ be defined as in Eq. (21), i.e., $Q_\lambda = O_\lambda^\dagger P O_\lambda$, with O_λ as defined in Eq. (18) and P being a Pauli string. We consider the infinite norm $\|Q_\lambda\|_\infty$ and from now on we will drop the subindex for simplicity. Let us write the following inequality for the norm of a product of operators: for A and B being invertible operators, it holds that

$$\|AB\| \geq \|A\| \frac{1}{\|B^{-1}\|}. \tag{D1}$$

Equation (D1) can be proven by using the submultiplicity of the norm, namely,

$$\|A\| = \|ABB^{-1}\| \leq \|AB\| \|B^{-1}\|. \tag{D2}$$

Let us first bound the norm of O_λ and (O_λ^\dagger) . With this and by means of Eq. (D1), we will be able to lower bound $\|Q_\lambda\|$.

Let us take the norm of O_λ and insert its definition from Eq. (18):

$$\begin{aligned} \|O_\lambda\| &= \left\| \prod_{n \in \mu(\lambda)} e^{-ik_{2,n}} \prod_{m \in \nu(\lambda)} e^{-\beta \kappa_{1,m}} \right\| \stackrel{\text{Eq. (D1)}}{\geq} \underbrace{\left\| \prod_{n \in \mu(\lambda)} e^{-ik_{2,n}} \right\|}_{=1} \prod_{m \in \nu(\lambda)} \|e^{\beta \kappa_{1,m}}\|^{-1} = \prod_{m \in \nu(\lambda)} \|e^{\beta \kappa_{1,m}}\|^{-1} \\ &\geq \prod_{m \in \nu(\lambda)} e^{-\beta} = e^{-\beta |\nu(\lambda)|}. \end{aligned}$$

Here we have used that $[\kappa_{1,r}, \kappa_{1,s}] = 0$, so we can just write the inverse $(\prod_{m \in \nu(\lambda)} e^{-\beta \kappa_{1,m}})^{-1} = \prod_{m \in \nu(\lambda)} (e^{-\beta \kappa_{1,m}})^{-1}$. We can now use that $(e^{-\beta \kappa_{1,m}})^{-1} = e^{\beta \kappa_{1,m}}$ and that $\|e^{\beta \kappa_{1,m}}\|_\infty \leq e^\beta$. With this, it follows that

$$\|O_\lambda\| = \Omega(e^{-\beta|\nu_\lambda|}). \quad (\text{D3})$$

Let us now study (O_λ^\dagger) :

$$\begin{aligned} \|(O_\lambda^\dagger)^{-1}\| &= \left\| \left(\prod_{m \in \nu(\lambda)} e^{-\beta \kappa_{1,m}} \prod_{n \in \mu(\lambda)} e^{i \kappa_{2,n}} \right)^{-1} \right\| = \left\| \prod_{n \in \mu(\lambda)} e^{-i \kappa_{2,n}} \prod_{m \in \nu(\lambda)} e^{\beta \kappa_{1,m}} \right\| \leq \underbrace{\left\| \prod_{n \in \mu(\lambda)} e^{-i \kappa_{2,n}} \right\|}_{\text{subm.}} \prod_{m \in \nu(\lambda)} \|e^{\beta \kappa_{1,m}}\| \\ &\leq \prod_{m \in \nu(\lambda)} e^\beta = e^{\beta|\nu(\lambda)|}. \end{aligned}$$

The derivation above implies the following bound:

$$\frac{1}{\|(O_\lambda^\dagger)^{-1}\|} \geq e^{-\beta|\nu(\lambda)|}. \quad (\text{D4})$$

Note that in all the bounds above it is present the cardinality of $|\nu(\lambda)|$, so let us estimate this value in terms of the size $|\lambda|$:

$$|\nu(\lambda)| = \left| \bigcup_{j \in \lambda} \nu_j \right| \leq \sum_{j \in \lambda} |\nu_j| = O(|\lambda|z^{2r_1}), \quad (\text{D5})$$

where we have used that $|\nu_j| = O(z^{2r_1})$, where z is the degree of incidence of the graph at which vertex the particles are placed and r_1 is the radius of the operators κ_1 .

With all the derivations above, now we can get a lower bound on $\|Q_\lambda\|$:

$$\begin{aligned} \|Q_\lambda\| &= \|O_\lambda P O_\lambda^\dagger\| \underset{\text{Eq. (D1)}}{\geq} \|O_\lambda\| \|(P O_\lambda^\dagger)^{-1}\|^{-1} \\ &= \|O_\lambda\| \frac{1}{\|(O_\lambda^\dagger)^{-1} P\|}, \end{aligned}$$

where we have use that $P^{-1} = P$. Now we can use submultiplicity of the norm and Eq. (D4) to get

$$\frac{1}{\|(O_\lambda^\dagger)^{-1} P\|} \geq \frac{1}{\|(O_\lambda^\dagger)^{-1}\|} \underbrace{\frac{1}{\|P\|}}_{=1} \geq e^{-\beta|\nu(\lambda)|}. \quad (\text{D6})$$

With this and Eqs. (D3) and (D5) now we get

$$\boxed{\|Q_\lambda\| \geq e^{-2\beta|\nu(\lambda)|} = \Omega(e^{-2\beta|\lambda|z^{2r_1}})}. \quad (\text{D7})$$

This means that $\|Q_\lambda\|$ decays *at most* as $O(e^{-2\beta|\lambda|z^{2r_1}})$. If we consider now the terms such that $|\lambda|$ is constant [or at most goes as $O(\log N)$, with N being the system size], then we get that $\|Q_\lambda\|$ decays at most polynomially with the system size.

APPENDIX E: ANALYSIS OF THE VERIFICATION TEST

In this Appendix, we analyze the verification test for the setup presented in Sec. IV A, that is, in which the verifier can perform local measurements to the state that the prover has prepared. This case is analogous to the one presented in [28]. We modify it here by imposing further restrictions on the prover: we assume that he can only perform Pauli

measurements, which might be more suitable for experiments carried out with NISQ devices.

We first analyze the probability that, for a given tolerance threshold ϵ , the empirical mean $\langle h_j \rangle_P$ computed through the measurements (where the subindex P emphasizes that the verifier measures the state that the prover has prepared) is above such threshold, due to statistical noise produced by finite sampling. This statistical fluctuation depends on the variance of the distribution $\langle h_j \rangle_P$, which is computed by decomposing the operators in the Pauli basis. This variance is bounded in Appendix E 2.

1. Bounds on the fidelity and the number of measurement rounds

Let us denote by $\langle h_j \rangle_P$ the random variable result of estimating the expectation value of h_j by measuring the state in the Pauli basis, for a number of samples L_j . We define this random variable as $\langle h_j \rangle_P = \frac{1}{L_j} \sum_{k=1}^{L_j} \bar{h}_j^{(k)}$, where \bar{h}_j refers to the random variable associated to $\langle h_j \rangle_P$ computed with a single round [see next section and Eq. (E10) for a formal definition of \bar{h}]. As before, we assume that $L_j = L_j'$ and that we do not recycle samples, i.e., we always use different samples for different local terms h_j . In the same way, we denote by $\langle H \rangle_P$ the estimator of the energy, which is a random variable whose mean we denote by $\mathbb{E}[\langle H \rangle_P]$. We also define $F_P = 1 - \frac{\langle H \rangle_P}{\delta}$, with δ being the lower bound on the spectral gap computed in Sec. III. By Eq. (26), F_P is an estimator of a lower bound on the fidelity between the prover's state and the true ground state. Then, we get the following result: for $\alpha \in (0, 1)$ and a threshold parameter $\epsilon > 0$, it holds that

$$P(F_P < \mathbb{E}[F_P] - \epsilon) \leq \alpha, \quad (\text{E1})$$

as long as the number of samples L_j satisfies

$$L_j \geq \frac{2^{5|\lambda|+1}}{\delta^2 \epsilon^2} \ln \left(\frac{1}{\alpha^{1/N}} \right), \quad (\text{E2})$$

where N is the number of terms in the Hamiltonian and $|\lambda|$ is its locality.

Proof. By Eq. (26) we get that

$$P(F_P < \mathbb{E}[F_P] - \epsilon) = P(\langle H \rangle_P > \mathbb{E}[\langle H \rangle_P] + \delta \epsilon). \quad (\text{E3})$$

Recall that $\langle H \rangle_P = \sum_j \langle h_j \rangle_P$. We can assume that the numbers of samples used to estimate any $\langle h_j \rangle_P$ are large enough so

the central limit theorem applies. Thus, we can assume that $\langle h_j \rangle_P \sim \mathcal{N}(\mathbb{E}[\langle h_j \rangle_P], \text{Var}[\langle h_j \rangle_P])$ follows a Gaussian distribution, which therefore implies that $\langle H \rangle_P$ is Gaussian itself. Then, we can bound (E3) by

$$P(\langle H \rangle_P > \mathbb{E}[\langle H \rangle_P] + \delta\epsilon) \leq e^{-(\delta\epsilon)^2/2 \text{Var}[\langle H \rangle_P]} = \prod_j e^{-(\delta\epsilon)^2/2 \text{Var}[\langle h_j \rangle_P]}, \quad (\text{E4})$$

where the inequality follows from bounding the tails of the Gaussian distribution and the last equality comes from the fact that $\text{Var}[\langle H \rangle_P] = \sum_j \text{Var}[\langle h_j \rangle_P]$ since the random variables $\langle h \rangle_P$ are independent. Note that $\text{Var}[\langle h_j \rangle_P] = \text{Var}[\bar{h}_j]/L_j$. We will see in the following section that

$$\text{Var}[\bar{h}_j] \leq 2^{5|\lambda|}. \quad (\text{E5})$$

By substituting this into the previous equation and imposing the bound (E2) on L_j we get

$$P(\langle H \rangle_P > \mathbb{E}[\langle H \rangle_P] + \delta\epsilon) \leq \alpha, \quad (\text{E6})$$

which thus proves Eq. (E1). Note that this equation is similar to Eq. (7) and (15) in [28] and a derivation in terms of Hoeffding bounds as in this paper would also apply here, although the difference is again that we do not measure the eigenenergies of h directly but only the expectation values of strings of Pauli operators.

2. Bound on the variance of $\langle h \rangle_P$

Let us now study the properties of the distribution of $\langle h \rangle_P$ for a fixed local term h . We first introduce some preliminary definitions, some of which were already introduced in the main text, that will help in the upcoming discussion:

(i) We define a set random variables α_j , for $j = 1, \dots, N$. Each variable α_j can take the values $\{x, y, z\}$ with equal probability. We denote the set of this variables by B , $B = \{\alpha_j\}_{j=1}^N$. Note that each element of B corresponds to the basis in which we choose to measure each qubit independently.

(ii) We denote by J any of the subsets of indices of $\lambda(h) = \{j_1, \dots, j_{|\lambda|}\}$. Note that the cardinal of J goes from $0, \dots, |\lambda|$,

where we also include the empty set for convenience. In the same fashion, we denote by B_J the subset of random variables α_j as defined before with corresponding indices $j \in J$.

(iii) Let $s_{B_J}^J = s_{\alpha_{j_1}}^{j_1} \dots s_{\alpha_{j_{|J|}}}^{j_{|J|}}$ be the random variable with possible outputs ± 1 resulting from measuring the qubits $j \in J$ in the basis $\alpha_j \in B_J$.

With these preliminaries we can now define the following random variable \bar{h} :

$$\begin{aligned} \bar{h} = o_0 + \sum_{j=1}^{|\lambda|} 3 o_{\alpha_j}^j s_{\alpha_j}^j + \sum_{j,k=1}^{|\lambda|} 3^2 o_{\alpha_j, \alpha_k}^{j,k} s_{\alpha_j}^j s_{\alpha_k}^k + \dots \\ + 3^{|\lambda|} o_{\alpha_{j_1}^{j_1} \dots \alpha_{j_{|\lambda|}}^{j_{|\lambda|}}} s_{\alpha_{j_1}}^{j_1} \dots s_{\alpha_{j_{|\lambda|}}}^{j_{|\lambda|}} = \sum_{J \subseteq \lambda(h)} 3^{|J|} o_{B_J}^J s_{B_J}^J. \end{aligned} \quad (\text{E7})$$

Here the coefficients $o_{\alpha_{j_1}^{j_1} \dots \alpha_{j_{|\lambda|}}^{j_{|\lambda|}}}$ are those of the Pauli expansion of $\langle h \rangle$:

$$\begin{aligned} \langle h \rangle = \sum_{j=1}^n \sum_{\gamma=x,y,z} o_{\gamma}^j \langle \sigma_{\gamma}^j \rangle + \sum_{j,k=1}^n \sum_{\gamma_1, \gamma_2=x,y,z} o_{\gamma_1 \gamma_2}^{j,k} \langle \sigma_{\gamma_1}^j \sigma_{\gamma_2}^k \rangle + \dots \\ + \sum_{j_1, \dots, j_k=1}^n \sum_{\gamma_{j_1}, \dots, \gamma_{j_k}=x,y,z} o_{\gamma_{j_1}^{j_1} \dots \gamma_{j_k}^{j_k}} \langle \sigma_{\gamma_{j_1}}^{j_1} \dots \sigma_{\gamma_{j_k}}^{j_k} \rangle. \end{aligned} \quad (\text{E8})$$

Note that in the definition of the variable \bar{h} in Eq. (E7) we do not sum over all possible basis choices: this variable actually corresponds to the computation of $\langle h \rangle_P$ with a single round, in which we only measure the state once in some basis $\alpha = (\alpha_1, \dots, \alpha_N)$. Note that this is different from Eq. (30), in which we averaged over different rounds with the same basis choice [one could in principle also perform an analysis with multirounds, but then $\mathcal{R}(\boldsymbol{\gamma})$ is itself a random variable, which makes the analysis more cumbersome]. With this single sample, we can only give an estimator for the following marginals:

$$\langle \sigma_{\alpha_j}^j \rangle, \langle \sigma_{\alpha_{j_1}}^{j_1} \sigma_{\alpha_{j_2}}^{j_2} \rangle, \dots, \langle \sigma_{\alpha_{j_1}}^{j_1} \dots \sigma_{\alpha_{j_{|\lambda|}}}^{j_{|\lambda|}} \rangle, \quad (\text{E9})$$

where $\alpha_{j_k} \in B$. Let us now compute the expectation value of \bar{h} :

$$\begin{aligned} \mathbb{E}[\bar{h}] = o_0 + \sum_{j=1}^{|\lambda|} \sum_{\gamma_j=x,y,z} 3P(\alpha_j = \gamma_j) \mathbb{E}[o_{\alpha_j}^j s_{\alpha_j}^j] + \sum_{j,k=1}^{|\lambda|} \sum_{\gamma_j, \gamma_k=x,y,z} 3^2 P(\alpha_j = \gamma_j, \alpha_k = \gamma_k) \mathbb{E}[o_{\alpha_j, \alpha_k}^{j,k} s_{\alpha_j}^j s_{\alpha_k}^k] \\ + \dots + \sum_{\gamma_{j_1}, \dots, \gamma_{j_{|\lambda|}}} 3^{|\lambda|} P(\alpha_{j_1} = \gamma_{j_1}, \dots, \alpha_{j_{|\lambda|}} = \gamma_{j_{|\lambda|}}) \mathbb{E}[o_{\alpha_{j_1}^{j_1} \dots \alpha_{j_{|\lambda|}}^{j_{|\lambda|}}} s_{\alpha_{j_1}}^{j_1} \dots s_{\alpha_{j_{|\lambda|}}}^{j_{|\lambda|}}] \\ = \sum_{J \subseteq \lambda(h)} \sum_{\gamma_{j_1}, \dots, \gamma_{j_{|\lambda|}}=x,y,z} 3^{|J|} P(B_J = (\gamma_{j_1}, \dots, \gamma_{j_{|\lambda|}})) o_{B_J}^J \mathbb{E}[s_{B_J}^J]. \end{aligned} \quad (\text{E10})$$

Since α_j can take the values $\{x, y, z\}$ with equal probability we have that $P(\alpha_j = \gamma_j) = \frac{1}{3}, \forall j, \forall \gamma_j$. Moreover, the variables α_j are independent and therefore the joint probability can be factorized:

$$P(\alpha_{j_1} = \gamma_{j_1}, \dots, \alpha_{j_k} = \gamma_{j_k}) = P(\alpha_{j_1} = \gamma_{j_1}) \dots P(\alpha_{j_k} = \gamma_{j_k}) = \left(\frac{1}{3}\right)^k,$$

for $k = 1, \dots, |\lambda|$. Note that Eq. (E10) leads to the same expression as Eq. (E8) and, therefore,

$$\langle h \rangle = \mathbb{E}[\bar{h}]. \quad (\text{E11})$$

In statistical terms, \bar{h} is an unbiased estimator for $\langle h \rangle$. We can now compute the variance of the variable \bar{h} . Now, note that $s_{B_j}^J$ and $s_{B_{j'}}^{J'}$ are not independent since J and J' may have nonempty intersection. Therefore, we express the variance of \bar{h} , defined by the sum in Eq. (E7), in terms of the covariance of its terms:

$$\text{Var}[\bar{h}] = \sum_{|J|,|J'|=0}^{|\lambda|} \sum_{J,J' \subseteq \lambda(h)} 3^{|J|} 3^{|J'|} a_{B_j}^J a_{B_{j'}}^{J'} \text{Cov}[s_{B_j}^J, s_{B_{j'}}^{J'}]. \tag{E12}$$

We can bound the coefficients $a_{B_j}^J a_{B_{j'}}^{J'}$ by $(o_{\max})^2$, where $o_{\max}^2 = \max\{(a_{B_j}^J)^2\}$ and take it out of the sum for simplicity. Note that the covariance that appears in (E12) can be bounded by

$$\text{Cov}[s_{B_j}^J, s_{B_{j'}}^{J'}] \leq 1. \tag{E13}$$

Then,

$$\text{Var}[\bar{h}] \leq (o_{\max})^2 \sum_{|J|,|J'|=0}^{|\lambda|} 3^{|J|} 3^{|J'|} \sum_{J,J' \subseteq \lambda(h)} 1 = (o_{\max})^2 \left(\sum_{|J|=1}^{|\lambda|} 3^{|J|} \sum_{J \subseteq \lambda(h)} 1 \right)^2 = (o_{\max})^2 \underbrace{\left(\sum_{|J|=1}^{|\lambda|} \binom{|\lambda|}{|J|} 3^{|J|} \right)^2}_{4^{|\lambda|}} = o_{\max}^2 4^{2|\lambda|}. \tag{E14}$$

Note that o_{\max}^2 depends on the particular local term that we are considering. We can get rid of this dependence by bounding it by the following value:

$$o_{\max}^2 \leq 2^{|\lambda|}, \tag{E15}$$

which follows from the fact that $\text{tr}(h^2) \leq 2^{|\lambda|}$.

[1] C. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland *et al.*, Experimental entanglement of four particle, *Nature (London)* **404**, 256 (2000).

[2] D. Leibfried, E. Knill, S. Seidelin, J. Britton, R. B. Blakestad, J. Chiaverini, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer *et al.*, Creation of a six-atom ‘Schrödinger cat’ state, *Nature (London)* **438**, 639 (2005).

[3] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harl, W. Hänsel, M. Hennrich, and R. Blatt, 14-Qubit Entanglement: Creation and Coherence, *Phys. Rev. Lett.* **106**, 130506 (2011).

[4] N. Friis, O. Marty, C. Maier, C. Hempel, M. Holzäpfel, P. Jurcevic, M. B. Plenio, M. Huber, C. Roos, R. Blatt *et al.*, Observation of Entangled States of a Fully Controlled 20-Qubit System, *Phys. Rev. X* **8**, 021012 (2018).

[5] I. Pogorelov, T. Feldker, Ch. D. Marciniak, L. Postler, G. Jacob, O. Kriegelsteiner, V. Podlesnic, M. Meth, V. Negnevitsky, M. Stadler *et al.*, A compact ion-trap quantum computing demonstrator, *PRX Quantum* **2**, 020343 (2021).

[6] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. Wang, S. Ebadi, H. Bernien, A. Zibrov, H. Pichler, S. Choi *et al.*, Generation and manipulation of Schrödinger cat states in Rydberg atom arrays, *Science* **365**, 570 (2019).

[7] L. DiCarlo, M. D. Reed, L. Sun, B. R. Johnson, J. M. Chow, J. M. Gambetta, L. Frunzio, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, Preparation and measurement of three-qubit entanglement in a superconducting circuit, *Nature (London)* **467**, 574 (2010).

[8] C. Song, K. Xu, W. Liu, C. Yang, S. Zheng, H. Deng, Q. Xie, K. Huang, Q. Guo, L. Zhang *et al.*, 10-Qubit Entanglement and Parallel Logic Operations with a Superconducting Circuit, *Phys. Rev. Lett.* **119**, 180511 (2017).

[9] M. Gong, M. Chen, Y. Zheng, S. Wang, C. Zha, H. Deng, Z. Yan, H. Rong, Y. Wu, S. Li *et al.*, Genuine 12-Qubit Entanglement on a Superconducting Quantum Processor, *Phys. Rev. Lett.* **122**, 110501 (2019).

[10] C. Song, K. Xu, H. Li, Y.R. Zhang, X. Zhang, W. Liu, Q. Guo, Z. Wang, W. Ren, J. Hao *et al.*, Generation of multicomponent atomic Schrödinger cat states of up to 20 qubits, *Science* **365**, 574 (2019).

[11] D. Bouwmeester, J.-W. Pan, M. Daniell, H. Weinfurter, and A. Zeilinger, Observation of Three-Photon Greenberger-Horne-Zeilinger Entanglement, *Phys. Rev. Lett.* **82**, 1345 (1999).

[12] J.-W. Pan, M. Daniell, S. Gasparoni, G. Weihs, and A. Zeilinger, Experimental Demonstration of Four-Photon Entanglement and High-Fidelity Teleportation, *Phys. Rev. Lett.* **86**, 4435 (2001).

[13] H.-S. Zhong, Y. Li, Wei Li, L.-C. Peng, Z.-E. Su, Y. Hu, Y.-M. He, X. Ding, W. Zhang, H. Li *et al.*, 12-Photon Entanglement and Scalable Scattershot Boson Sampling with Optimal Entangled-Photon Pairs from Parametric Down-Conversion, *Phys. Rev. Lett.* **121**, 250505 (2018).

[14] X.-L. Wang, Y.-H. Luo, H.-L. Huang, M.-C. Chen, Z.-E. Su, C. Liu, C. Chen, W. Li, Y.-Q. Fang, X. Jiang *et al.*, 18-Qubit Entanglement with Six Photons’ Three Degrees of Freedom, *Phys. Rev. Lett.* **120**, 260502 (2018).

[15] R. Laflamme, E. Knill, W. Zurek, P. Catasti, and S. Mariappan, NMR Greenberger–Horne–Zeilinger states, *Philos. Trans. R. Soc. London A* **356**, 1941 (1998).

[16] P. Neumann, N. Mizuochi, F. Rempp, P. Hemmer, H. Watanabe, S. Yamasaki, V. Jacques, T. Gaebel, F. Jelezko, and J. Wrachtrup, Multipartite entanglement among single spins in diamond, *Science* **320**, 1326 (2008).

[17] I. Cirac, D. Perez-Garcia, N. Schuch, and F. Verstraete, Matrix product states and projected entangled pair states: concepts, symmetries, and theorems, *Rev. Mod. Phys.* **93**, 045003 (2021).

- [18] M. Fannes, B. Nachtergaele, and R. F. Werner, Finitely correlated states on quantum spin chains, *Commun. Math. Phys.* **144**, 443 (1992).
- [19] F. Verstraete and J. I. Cirac, Valence-bond states for quantum computation, *Phys. Rev. A* **70**, 060302(R) (2004).
- [20] A. Broadbent, J. Fitzsimons, and E. Kashefi, *Universal Blind Quantum Computation*, *Proceedings of the Annual IEEE Symposium on Foundations of Computer Science, FOCS* (IEEE, Piscataway, NJ, 2008).
- [21] B. W. Reichardt, F. Unger, and U. Vazirani, A Classical Leash for a Quantum System: Command of Quantum Systems via Rigidity of CHSH Games, in *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science, ITCS '13 2013* (Association for Computing Machinery, New York, 2013), p. 321–322.
- [22] B. W. Reichardt, F. Unger, and U. Vazirani, Classical command of quantum systems, *Nature (London)* **496**, 456 (2013).
- [23] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, Hamiltonian Learning and Certification Using Quantum Resources, *Phys. Rev. Lett.* **112**, 190501 (2014).
- [24] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, Quantum Hamiltonian learning using imperfect quantum resources, *Phys. Rev. A* **89**, 042314 (2014).
- [25] N. Wiebe, C. Granade, and D. G. Cory, Quantum bootstrapping via compressed quantum Hamiltonian learning, *New J. Phys.* **17**, 022005 (2015).
- [26] D. Aharonov, M. Ben-Or, E. Eban, and U. Mahadev, Interactive proofs for quantum computations, [arXiv:1704.04487](https://arxiv.org/abs/1704.04487).
- [27] J. F. Fitzsimons, M. Hajdušek, and T. Morimae, *Post hoc* Verification of Quantum Computation, *Phys. Rev. Lett.* **120**, 040501 (2018).
- [28] D. Hangleiter, M. Kliesch, M. Schwarz, and J. Eisert, Direct certification of a class of quantum simulations, *Quantum Sci. Technol.* **2**, 015004 (2017).
- [29] U. Mahadev, Classical Verification of Quantum Computations, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, Piscataway, NJ, 2018), pp. 259–267.
- [30] Z. Brakerski, P. Christiano, U. Mahadev, U. Vazirani, and T. Vidick, A Cryptographic Test of Quantumness and Certifiable Randomness from a Single Quantum Device, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, Piscataway, NJ, 2018), pp. 320–331.
- [31] S. Aaronson and L. Chen, *Complexity-Theoretic Foundations of Quantum Supremacy Experiments*, in *Proceedings of the 32nd Computational Complexity Conference, CCC '17*, Dagstuhl, DEU (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017).
- [32] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [33] C. Schön, K. Hammerer, M. M. Wolf, J. I. Cirac, and E. Solano, Sequential generation of matrix-product states in cavity QED, *Phys. Rev. A* **75**, 032311 (2007).
- [34] J.-C. Besse, K. Reuer, M. C. Collodo, A. Wulff, L. Wernli, A. Copetudo, D. Malz, P. Magnard, A. Akin, M. Gabureac *et al.*, Realizing a deterministic source of multipartite-entangled photonic qubits, *Nat. Commun.* **11**, 4877 (2020).
- [35] I. Schwartz, D. Cogan, E. R. Schmidgall, Y. Don, L. Gantz, O. Kenneth, N. H. Lindner, and D. Gershoni, Deterministic generation of a cluster state of entangled photons, *Science* **354**, 434 (2016).
- [36] D. Istrati, Y. Pilnyak, J. C. Loredó, C. Antn, N. Somaschi, P. Hilaire, H. Ollivier, M. Esmann, L. Cohen, L. Vidro *et al.*, Sequential generation of linear cluster states from a single photon emitter, *Nat. Commun.* **11**, 5501 (2020).
- [37] S. Takeda, K. Takase, and A. Furusawa, On-demand photonic entanglement synthesizer, *Sci. Adv.* **5**, eaaw4530 (2019).
- [38] M. C. Bañuls, D. Pérez-García, M. M. Wolf, F. Verstraete, and J. I. Cirac, Sequentially generated states for the study of two-dimensional systems, *Phys. Rev. A* **77**, 052306 (2008).
- [39] H. Pichler, S. Choi, P. Zoller, and M. D. Lukin, Universal photonic quantum computation via time-delayed feedback, *Proc. Natl. Acad. Sci. USA* **114**, 11362 (2017).
- [40] M. Gimeno-Segovia, T. Rudolph, and S. E. Economou, Deterministic Generation of Large-Scale Entangled Photonic Cluster State from Interacting Solid State Emitters, *Phys. Rev. Lett.* **123**, 070501 (2019).
- [41] Y. Ge, A. Molnár, and J. I. Cirac, Rapid Adiabatic Preparation of Injective Projected Entangled Pair States and Gibbs States, *Phys. Rev. Lett.* **116**, 080503 (2016).
- [42] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Computational Complexity of Projected Entangled Pair States, *Phys. Rev. Lett.* **98**, 140506 (2007).
- [43] J. Haferkamp, D. Hangleiter, J. Eisert, and M. Gluza, Contracting projected entangled pair states is average-case hard, *Phys. Rev. Res.* **2**, 013010 (2020).
- [44] S. Jansen, M.-B. Ruskai, and R. Seiler, Bounds for the adiabatic approximation with applications to quantum computation, *J. Math. Phys.* **48**, 102111 (2007).
- [45] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [46] E. H. Lieb and D. W. Robinson, The finite group velocity of quantum spin systems, *Commun. Math. Phys.* **28**, 251 (1972).
- [47] M. Sanz, D. Perez-Garcia, M. M. Wolf, and J. I. Cirac, A quantum version of Wielandts inequality, *IEEE Trans. Inf. Theory* **56**, 4668 (2010).
- [48] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Peps as Unique Ground States of Local Hamiltonians, *Quantum Inf. Comput.* **8**, 650 (2008).
- [49] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Criticality, the Area Law, and the Computational Power of Projected Entangled Pair States, *Phys. Rev. Lett.* **96**, 220601 (2006).
- [50] M. J. Kastoryano and A. Lucia, Divide and conquer method for proving gaps of frustration free Hamiltonians, *J. Stat. Mech.* (2018) 033105.
- [51] J. Kempe, A. Kitaev, and O. Regev, The Complexity of the Local Hamiltonian Problem, *SIAM J. Comput.* **35**, 1070 (2006).
- [52] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, Efficient quantum state tomography, *Nat. Commun.* **1**, 149 (2010).
- [53] J. Cotler and F. Wilczek, Quantum Overlapping Tomography, *Phys. Rev. Lett.* **124**, 100401 (2020).
- [54] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, R. G. S. L. Brandao, D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).

- [55] R. Raussendorf and H. J. Briegel, Quantum computing via measurements only, [arXiv:quant-ph/0010033](https://arxiv.org/abs/quant-ph/0010033).
- [56] J. Miller and A. Miyake, Hierarchy of universal entanglement in 2D measurement-based quantum computation, [npj Quantum Inf.](https://doi.org/10.1038/npj-quantum-inf.2016.2) **2**, 16036 (2016).
- [57] J. Bermejo-Vega, D. Hangleiter, M. Schwarz, R. Raussendorf, and J. Eisert, Architectures for Quantum Simulation Showing a Quantum Speedup, [Phys. Rev. X](https://doi.org/10.1103/PhysRevX.8.021010) **8**, 021010 (2018).
- [58] M. J. Bremner, R. Jozsa, and D. J. Shepherd, Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy, [Proc. R. Soc. A: Math., Phys. Eng. Sci.](https://doi.org/10.1093/imrn/2011/467) **467**, 459 (2011).
- [59] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations, [Phys. Rev. Lett.](https://doi.org/10.1103/PhysRevLett.117.080501) **117**, 080501 (2016).
- [60] I. Šupić and J. Bowles, Self-testing of quantum systems: A review, [Quantum](https://doi.org/10.22331/q-2020-04-337) **4**, 337 (2020).
- [61] Y.-K. Liu, Consistency of Local Density Matrices Is QMA-Complete, in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, edited by J. Díaz, K. Jansen, J. D. P. Rolim, and U. Zwick (Springer, Berlin, 2006), pp. 438–449.
- [62] A. Broadbent and A. B. Grilo, QMA-hardness of Consistency of Local Density Matrices with Applications to Quantum Zero-Knowledge, in *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, Piscataway, NJ, 2020).
- [63] T. J. Osborne, Simulating adiabatic evolution of gapped spin systems, [Phys. Rev. A](https://doi.org/10.1103/PhysRevA.75.032321) **75**, 032321 (2007).
- [64] S. Bravyi and B. Terhal, Complexity of stoquastic frustration-free Hamiltonians, [SIAM J. Comput.](https://doi.org/10.1137/S0895472109371462) **39**, 1462 (2010).
- [65] H. J. Briegel and R. Raussendorf, Persistent Entanglement in Arrays of Interacting Particles, [Phys. Rev. Lett.](https://doi.org/10.1103/PhysRevLett.86.910) **86**, 910 (2001).
- [66] M. Hein, J. Eisert, and H. J. Briegel, Multipartite entanglement in graph states, [Phys. Rev. A](https://doi.org/10.1103/PhysRevA.69.062311) **69**, 062311 (2004).
- [67] L. Hartmann, J. Calsamiglia, W. Dür, and H. Briegel, Weighted graph states and applications to spin chains, lattices and gases, [J. Phys. B: At., Mol. Opt. Phys.](https://doi.org/10.1088/1464-4266/40/1/S1) **40**, S1 (2007).
- [68] A. Kitaev, Fault-tolerant quantum computation by anyons, [Ann. Phys.](https://doi.org/10.1088/1126-6708/2003/02/038) **303**, 2 (2003).
- [69] S. Bravyi and M. Vyalyi, Commutative version of the local Hamiltonian problem and common eigenspace problem, [Quantum Inf. Comput.](https://doi.org/10.1007/s10003-005-0005-0) **5**, 187 (2005).
- [70] D. Aharonov and L. Eldar, *On the Complexity of Commuting Local Hamiltonians, and Tight Conditions for Topological Order in Such Systems*, 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (IEEE, Piscataway, NJ, 2011), p. 334.