

Efficient algorithm for approximating one-dimensional ground states

Dorit Aharonov*

School of Computer Science and Engineering, Hebrew University, Jerusalem, Israel

Itai Arad†

School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel

Sandy Irani‡

Computer Science Department, University of California, Irvine, California, USA

(Received 17 November 2009; published 16 July 2010)

The density-matrix renormalization-group method is very effective at finding ground states of one-dimensional (1D) quantum systems in practice, but it is a heuristic method, and there is no known proof for when it works. In this article we describe an efficient classical algorithm which provably finds a good approximation of the ground state of 1D systems under well-defined conditions. More precisely, our algorithm finds a matrix product state of bond dimension D whose energy approximates the minimal energy such states can achieve. The running time is exponential in D , and so the algorithm can be considered tractable even for D , which is logarithmic in the size of the chain. The result also implies trivially that the ground state of any local commuting Hamiltonian in 1D can be approximated efficiently; we improve this to an exact algorithm.

DOI: [10.1103/PhysRevA.82.012315](https://doi.org/10.1103/PhysRevA.82.012315)

PACS number(s): 03.67.Ac, 05.50.+q, 89.70.Eg

I. INTRODUCTION

Finding ground states of local one-dimensional (1D) Hamiltonian systems is a major problem in physics. The most commonly used method is density matrix renormalization group (DMRG) [1–6], discovered in 1992. DMRG can be cast in the form of matrix product states (MPSs) which are succinct representations of 1D quantum states using $D \times D$ matrices, where the coefficients in the state can be written in terms of products of these matrices. The number of matrices is dn , where d is the dimension of each individual particle and n is the number of particles in the system. The parameter D is called the *bond dimension*. DMRG works essentially as follows: The algorithm starts with some initial MPS and sweeps from one end of the chain to the other, optimizing the entries of the matrices at one site with the other parameters fixed. Some versions allow optimizing over two neighboring sites at once, which enables the algorithm to increase the bond dimension in the course of the algorithm for improved accuracy. In all cases, the approach is to apply local optimizations iteratively. It is thus easy to construct examples in which the DMRG algorithm gets trapped in a local minimum. To illustrate this, think of a 1D spin chain whose Hamiltonian consists of two types of interactions: One type consists of interactions which force the spins to be aligned; every two neighboring sites gain an energy penalty of say 4 if they are not aligned. The other type of term gives every spin an energy penalty of 1 if it points upward. Starting from the all-up string, a local move only increases the energy; thus, local update rules cannot take the system to its ground state, the all-down string. This example can of course be handled by randomizing the initial string, for example, or increasing

the window size; however, it demonstrates that DMRG has a fundamental difficulty in addressing nonlocal characteristics of the system. It is natural to ask if there is a general algorithm that does not get stuck in local minima as DMRG does and provably always find a good approximation of the ground state of a given 1D system in a reasonable amount of time.

To answer this question, we first ask what is known regarding the analogous question in the easier, classical, case. It was Kitaev [7] who drew the important connection between the problem of finding ground energy and ground states of local Hamiltonians and the well-known classical *constraint satisfaction problem* (CSP). The input to a CSP consists of constraints $\{H_c\}_c$ on n q -state classical particles. Each H_c acts on k particles (for some constant k) and is given as a Boolean function on the possible assignments to those k particles; when $H_c = 1$ the configuration is forbidden and when $H_c = 0$ it is allowed. The problem is to determine the maximum number of constraints that can be satisfied or, alternatively, to minimize $\sum_c H_c$. The decision version of this problem is to determine whether it is possible to satisfy more than some given number of constraints. This is one of the most well-known NP-complete (where NP denotes nondeterministic polynomial) problems. CSP can clearly be seen as a special case of the problem of finding ground states and ground energies of local Hamiltonians, in which the terms in the Hamiltonian are projections on local forbidden configurations. This analogy has led over the past few years to many interesting insights regarding the local Hamiltonian problem (see, e.g., Refs. [7–12]).

Let us therefore see what the known classical results regarding CSP in 1D can teach us about 1D local Hamiltonians and their ground states. We recall that in the classical case, 1D CSPs (in which the particles are arranged in a line and constraints are between k adjacent neighbors) are dramatically easier than their higher-dimensional counterparts. While even the 2D case is NP complete, the 1D problem can be solved

*dorit.aharonov@gmail.com

†arad.itai@gmail.com

‡irani@ics.uci.edu

in polynomial time. The reason for the tractability of the problem in 1D is essentially that the problem can be divided into subproblems, namely, the left- and the right-hand sides of the chain, which interact only via the k particles on the border. The fact that these particles can only be assigned a small number of possible values makes it possible to handle the problem by solving each subproblem separately for each fixed possible assignment to the border particles and then gluing the subsolutions together by picking the best choice for the middle particles. We explain the algorithm in detail later; the outcome is an algorithm which is linear in the number of particles in the chain and quadratic in the number of states per particle.

Unfortunately, there is no hope of getting such a general result for the 1D quantum problem. Aharonov *et al.* [10] have shown that approximating the ground energy for general 1D quantum systems is as hard as quantum-NP. Even when restricted to ground states that are well-approximated by MPSs of polynomial bond dimension, the problem is NP-hard, as was shown by Schuch *et al.* [13]. A related earlier result due to Eisert [14] showed that optimizing a constant number of matrices in the MPS representation subject to fixed values in the other matrices is NP-hard. These results indicate that the dichotomy between the computational difficulty of 1D and 2D classical systems does not carry over to the quantum setting, and it is highly unlikely that the quantum 1D problem is tractable. Nevertheless, we show here that using the classical 1D algorithm as a template for an algorithm for the quantum problem leads to a solution for a wide and interesting class of local Hamiltonian problems, namely, for those cases in which we can assume that the bond dimension is small.

A. Main result

We derive an efficient algorithm for approximating the minimal energy of a 1D system among all states of a bounded bond dimension D . The algorithm is exponential in D and thus can be considered reasonable, though maybe not practical, even for D , which is logarithmic in the size of the chain. The algorithm also provides a description of an MPS with the approximate minimal energy.

Theorem 1. Let H be a nearest-neighbor Hamiltonian on a 1D system of n d -dimensional particles. Let J be a bound on the operator norm of each local term. There is an algorithm that takes as input ϵ , H , and D and produces an MPS $|\Omega\rangle$ of bond dimension D , such that for any MPS $|\psi\rangle$ of bond dimension D with $nD^2 \geq 12$,

$$\langle \Omega | H | \Omega \rangle \leq \langle \psi | H | \psi \rangle + 2JD^2n^2\epsilon. \quad (1)$$

The algorithm runs in time $n \cdot \text{poly}(d, D, N)$, where $N = O\left(\frac{144dD}{\epsilon}\right)^{D+2dD^2}$.

Several remarks are in place here. First, note that the restriction that the interactions are nearest neighbor is done without loss of generality since any 1D system can be reduced to a 2-local 1D system with nearest-neighbor interactions by grouping neighboring particles together.

Note also that the running time in the preceding theorem is phrased as a linear function in n , the size of the system, times some fixed amount of time spent per particle. The error, however, scales with n^2 . One may want to apply the theorem to derive an approximation with a fixed additive error δ , in

which case simply set $\epsilon = \delta n^{-2}$ in the preceding theorem to get the running time as a function of δ .

This result shows that the problem of finding bounded bond dimension MPSs can be done in polynomial time. Unfortunately, the running time, though efficient in theory, is quite impractical, as even for $D = 2$ and the error ϵ/n^2 a constant, we get a running time which scales like n^{16} . It is hard to imagine that these running times are practical. Nevertheless, it is very likely that the running time can be improved; in particular, when solving specific problems with certain symmetries, dramatic improvements may be possible. Moreover, it is possible that this algorithm can be used to boost DMRG in certain cases where it gets stuck or to create the initial state of DMRG. All these improvements are left for further research.

We now provide an overview of the algorithm. To understand the general idea, we first recall how the classical 1D algorithm works in detail. Consider the case of the classical CSP on a line with $k = 2$, namely, the problem of minimizing the energy function $H = \sum_{i=1}^n H_{i,i+1}$. An optimal assignment can be found efficiently by a standard algorithmic technique called *dynamic programming*. Define the partial problem up to the $(r+1)$ th particle, $H_r = \sum_{i=1}^r H_{i,i+1}$. The algorithm starts with the partial problem defined for $r = 1$ and creates a list L_2 of possible assignments to the first two particles as follows: For each of the q possible assignments σ_2 to particle 2, the algorithm finds an assignment σ_1 to particle 1 which minimizes $H_1(\sigma_1, \sigma_2)$. That optimal σ_1 is called the *tail* of σ_2 . For each σ_2 the algorithm keeps its tail σ_1 and also the energy of this partial assignment, $H_1(\sigma_1, \sigma_2)$. L_2 thus contains the best possible partial assignment with each possible ending. After $r - 1$ iterations, we assume the algorithm has a list L_r consisting of an optimal tail $\sigma_1, \dots, \sigma_{r-1}$ for each of the q possible assignments σ_r to the r th particle, where optimality is measured with respect to H_{r-1} . In other words, the algorithm has a solution to the subproblem confined to the first r particles, with any possible ending. To include the next particle, and create the next list L_{r+1} , the algorithm finds the optimal tail of each assignment σ_{r+1} . This is done by considering all items in the list L_r as possible tails for σ_{r+1} and taking the tail which minimizes $H_r(\sigma_1, \dots, \sigma_{r+1})$. In each of the $n - 1$ iterations, the algorithm checks for each of the q possible assignments σ_r , all q items in the list L_{r-1} . Thus, in time which is linear in n and quadratic in q , we can derive the final list L_{n-1} . The final solution is an assignment of minimal energy in that list.

The main idea in this article is to generalize the preceding algorithm to MPSs by replacing assignments to particles by possible values of MPS matrices. Since matrices are continuous objects, we use an ϵ net over all possible matrices of bond dimension D . The number of possible assignments to one variable, q , will now be replaced by the number of points in the ϵ net, denoted as N . We will move from one site to the next, keeping track of the minimum-energy MPS state, which ends in each MPS matrix for the rightmost particle that the algorithm has reached.

In order to carry out this idea, it must not happen that the choice of the MPS matrix of a later iteration can change the optimality of the partial MPS state found in an earlier iteration. To avoid this, we work with a restricted form of MPSs called *canonical MPSs*, in which the energy of each term in

the Hamiltonian depends only on MPS matrices associated with nearby particles. There are, however, various technical issues we need to handle. In particular, we cannot use perfectly canonical MPSs but only an approximated version of those, which imposes further technicalities, and in particular, the neighboring MPS matrices do not match perfectly (we call this *imperfect stitching*). These technicalities make the error analysis a bit subtle. Before we formally define canonical MPSs and provide the details of the algorithm, we mention an implication for a related problem.

B. Commuting Hamiltonians in 1D

A problem related to finding minimum-energy MPS states is the complexity of calculating the ground energy of commuting Hamiltonians in which all the local terms commute. Bravyi and Vyalyi proved that for 2-local Hamiltonians the problem lies inside NP [8]. For k -local commuting Hamiltonians with $k > 2$, the complexity of the problem is still open. The complexity of the 1D case was not studied before as far as we know; an immediate corollary of Theorem 1 is that there is an efficient classical algorithm for approximating the ground energy of commuting Hamiltonians in 1D to within $1/\text{poly}(n)$. This is because the ground state of a commuting Hamiltonian in 1D is an MPS of constant D (this is a well-known fact that we explain later for completeness), and therefore Theorem 1 can be applied. In fact, the result can be improved to an exact algorithm (up to exponentially good approximations due to truncations of real numbers) for a certain general class of problems. We prove the following.

Theorem 2. Given is a 1D Hamiltonian whose terms commute. There is an efficient algorithm that can compute the ground energy of this Hamiltonian to within any desired accuracy ϵ in time polynomial in n and in $\frac{1}{\epsilon}$. If we may assume also that the ground space of the total Hamiltonian is well separated from the higher excited states, by a spectral gap which is at least $1/\text{poly}(n)$, then the algorithm can find both the ground energy and a description of an MPS for the ground state exactly (i.e., up to exponentially small errors due to handling of real numbers).

The basic idea for the exact algorithm can be illustrated when the terms in the Hamiltonian are all projections and the ground state is unique. Since the terms commute, the ground state is an eigenstate of each term separately, with eigenvalue either 0 or 1. We start by applying the dynamic programming algorithm, to create a good approximation of the ground state. From this approximation we can deduce the correct eigenvalue (0 or 1) for each of the terms. The projections on the relevant eigenspaces can then be applied to the MPS of the approximate state to make it exact. One gets a tensor network of small depth, which can be converted into an MPS again. It can be shown that applying the projections does not increase the bond dimension of the MPS too much with respect to the approximating state. The details are fleshed out in the proof (Sec. V).

Handling the degenerate case is very easy; essentially, we force the dynamic algorithm to choose one state of the various possible states. The assumption on the spectral gap ensures that the errors created by the ϵ -net approximations would not cause a confusion between the ground space and some excited states.

We provide an alternative proof of Theorem 2, which also uses dynamic programming. In fact, this proof holds for a somewhat stronger version of the theorem, in which the conditions on the spectrum are far less restrictive. In the algorithm given by this approach, the state is not provided as an MPS but rather as a tensor product of two-particle states. The construction is based on the work of Ref. [8], in which it is proved that the ground states of 2-local commuting Hamiltonians have this special structure. Bravyi and Vyalyi use this structure to show that general 2-local commuting Hamiltonians are in NP. Since 1D chains with k -local interactions can always be made 2-local by treating nearby particles as one particle of a larger dimension, Ref. [8] implies that the 1D commuting problem lies in NP. However, by exploiting the special form of these ground states, dynamic programming can be applied to find the solution efficiently in a very similar manner to the 1D CSP, in which the NP witness is found using the 1D structure. Unfortunately, in this approach too, it seems that one cannot avoid some assumption on the spectrum of the total Hamiltonian, albeit a significantly less restrictive one. Throughout its execution, the dynamic programming algorithm compares various partial energies. If these are too close, and cannot be distinguished even by computations performed with exponentially good precision, then the algorithm might get confused between the ground energy and a slightly excited state. A sketch of the alternative proof of Theorem 2, providing the stronger version of it, and a discussion of the preceding precision issue are given in Sec. V.

We mention that this latter proof (and in particular the observation that dynamic programming can be useful for 1D quantum systems and not only for 1D classical systems) was the inspiration for the current article, rather than its corollary.

C. Discussion and open questions

It is natural to ask how much the results in this article can be improved. By Ref. [13], we know that no polynomial algorithm exists for finding optimal approximations of *polynomial* bond dimension (unless $P = NP$). However, the difficult instances of Ref. [13] have a spectral gap of $1/\text{poly}(n)$. Hastings has shown that ground states of 1D quantum systems with a constant gap can be approximated by MPSs with polynomial bond dimension [15]. However, this is too large to immediately yield an efficient algorithm from our result. It may still be true, however, that under the additional restriction that the Hamiltonian has a constant gap, a polynomial time algorithm exists, even when the bond dimension is as large as polynomial.

It is very likely that the efficiency of our algorithm can be significantly improved even for the general case. In particular, a factor of n would be shaved from the error in Theorem 1 if we could use an ϵ net which is both exactly canonical and enables perfect overlap between matrices at neighboring particles, as we later explain. Unfortunately, even if this can be done, the running time for this general algorithm is still quite large.

As mentioned earlier, we leave for further research the question of how this algorithm can be used in combination with DMRG and how certain symmetries in the problem can be utilized to enhance its performance time for specific interesting cases.




FIG. 1. MPS as a tensor network.

We note that very similar results to those presented in this article were derived independently by Schuch and Cirac [16].

D. Paper organization

Section II starts by defining tensor networks, MPSs, and canonical MPSs. In Sec. III we describe the algorithm. This is where the ϵ nets are defined and an algorithm to generate them is given. Also in Sec. III, we show how they are used in the dynamic programming algorithm. Section IV provides an exact analysis of the error accumulated in the algorithm. The complexity is analyzed as a function of the desired error. In Sec. V we provide the proof regarding the approximate and exact solutions for the commuting 1D case. We defer several technical lemmas to the Appendix.

II. TENSOR NETWORKS AND MATRIX PRODUCT STATES

A. Tensor networks

We start with some background on tensor networks, since MPSs are a special case of those. A detailed introduction to the use of tensor networks in the context of quantum computation can be found in Refs. [17–19].

A tensor network is a graph in which we allow some of the edges to be incident to only one node. These edges are called the *legs* of the network. Each node is assigned a tensor whose rank (number of indices) is equal to the degree of the node. Each index of the tensor corresponds to one edge that is incident to that node. To each edge (or index) we also assign a positive integer which indicates the range of the index. The indices associated with some of the edges in the tensor network may be assigned fixed values. The other edges are called *free* edges.

We call an assignment of values to the indices of the free edges in the network a *configuration*. With all the indices fixed, the tensor at each node in the network yields a particular value. We say that the value of the configuration is the product of the values for each of the nodes.

The value of the network is in general a tensor, whose rank is equal to the number of legs in the network. If there are no such legs, the value is simply a number (a scalar). Each assignment of values to the indices associated with the legs of the network gives rise to a value for the network tensor. We compute the tensor value for this assignment by summing over all configurations which are consistent with that assignment the value of each such configuration.

We note that often in the literature, one assigns values not to entire edges but to the two sides of an edge separately (where each side inherits its range of indices from the tensor associated with the node on that side). In the evaluation of the network, we require that the values on the two sides of one edge are equal, or else the entire configuration contributes zero to the sum.

Tensors will be denoted as bold-face fonts: λ, Γ, μ . Their contraction will be denoted as an expression like $\lambda \Gamma \mu$, when it is clear from the context along which indices the contraction is performed.

It is possible to restrict a tensor of rank k to a tensor of rank $k - 1$ by assigning a fixed value to one of its legs. For example, Γ_α is the restriction of the tensor Γ to the case in which the relevant edge associated with the index α is given some value (which, by the usual abuse of notation of variables and their values, will also be denoted as α).

It is convenient to associate with every tensor (which can be given as a contraction of a tensor network) a quantum state. For example, let $\Gamma = \Gamma_{\alpha,\beta}^i$ be a rank-3 tensor. Then we define $|\Gamma\rangle \stackrel{\text{def}}{=} \sum_{i,\alpha,\beta} \Gamma_{\alpha,\beta}^i |\alpha\rangle \otimes |\beta\rangle$.

B. Matrix product states

We work in the notation of Vidal [20] for MPSs, with minor changes. A MPS of a chain of n d -dimensional particles, with bond dimension D , is a tensor network with a 1D structure as in Fig. 1. Horizontal edges correspond to indices ranging from 1 to the bond dimension D and are denoted with α, β, \dots , while vertical edges correspond to indices ranging from 1 to the physical dimension d . (In our description, the end particles will actually have a different physical dimension, denoted d_{end} . This is required due to a technical reason described in Sec. II C.) The indices of vertical edges are denoted with i, j, \dots . The figures show two types of nodes: black and white. The tensors of black nodes are typically of rank 3 (except for the boundary tensors, which are of rank 2), and we denote them with Γ 's. For example, when the tensor that is second from left is written with its indices, it is denoted as $\Gamma_{\alpha_2, \alpha_3}^{[2]i}$, where the index [2] in brackets corresponds to its location in the graph. The tensors associated with white nodes are always of rank 2 and are denoted with λ 's. They are required to be diagonal and hence are given only one index (i.e., $\lambda_{\alpha_2}^{[2]i}$). Without loss of generality, we will also demand that the entries of λ are nonnegative since the phases can be absorbed in the neighboring Γ tensors.

The MPS defined by this network is $|\psi\rangle = \sum_{i_1, \dots, i_n} C_{i_1, \dots, i_n} |i_1\rangle \cdots |i_n\rangle$, with

$$C_{i_1, \dots, i_n} \stackrel{\text{def}}{=} \sum_{\alpha_2, \dots, \alpha_n} \Gamma_{\alpha_2}^{[1]i_1} \lambda_{\alpha_2}^{[2]i_2} \Gamma_{\alpha_2 \alpha_3}^{[2]i_2} \lambda_{\alpha_3}^{[3]i_3} \cdots \lambda_{\alpha_n}^{[n]i_n} \Gamma_{\alpha_n}^{[n]i_n}.$$

In the language of tensor states, $|\psi\rangle$ is exactly the tensor state of the contraction $\Gamma^{[1]i_1} \lambda^{[2]i_2} \Gamma^{[2]i_2} \cdots \lambda^{[n]i_n} \Gamma^{[n]i_n}$.

C. Canonical MPSs

An MPS is in *canonical form* if every cut in the chain induces a Schmidt decomposition (as in Fig. 2). In other words, we can rewrite the MPS by changing the order of summation to sum last over the index β of the j th λ tensor: $|\psi\rangle = \sum_{\beta} \lambda_{\beta}^{[j]i_j} |\mathbf{L}_{\beta}^{[j]i_j}\rangle \otimes |\mathbf{R}_{\beta}^{[j]i_j}\rangle$, where $\mathbf{L}_{\beta}^{[j]i_j}$ ($\mathbf{R}_{\beta}^{[j]i_j}$) denote the




FIG. 2. (Color online) A description of a canonical MPS. The tensors are chosen such that cutting a MPS between the $j - 1$ th and j th particles corresponds to the Schmidt decomposition between the left and right parts: $|\psi\rangle = \sum_{\beta} \lambda_{\beta}^{[j]} |\mathbf{L}_{\beta}^{[j]}\rangle \otimes |\mathbf{R}_{\beta}^{[j]}\rangle$.

contraction of the all the tensors to the left (right) of the cut with fixed β and $|\mathbf{L}_{\beta}^{[j]}\rangle$ ($|\mathbf{R}_{\beta}^{[j]}\rangle$) are their corresponding states. Then the canonical conditions are that for all j from 2 to n , $\sum_{\beta} (\lambda_{\beta}^{[j]})^2 = 1$ and $\langle \mathbf{L}_{\alpha}^{[j]} | \mathbf{L}_{\beta}^{[j]} \rangle = \langle \mathbf{R}_{\alpha}^{[j]} | \mathbf{R}_{\beta}^{[j]} \rangle = \delta_{\alpha\beta}$. In addition, for normalization, we require that the entire MPS state is normalized, which is guaranteed by the normalization requirement on the $\lambda^{[j]}$ tensors.

There is a small technical issue that needs attention: The canonical conditions cannot be satisfied at the boundaries if $d < D$. Consider for example the left boundary; there are not enough dimensions in the Hilbert space of the left particle for an orthonormal set of vectors $|\mathbf{L}_{\alpha}^{[2]}\rangle$ to exist. This issue remains a problem even as we move away from the boundary by one particle, as the dimension of the left-side Hilbert space increases to d^2 which may still be smaller than D . There are many ways of handling this technicality; here we choose to assume that the particles at the end of the chain have dimension of at least D . This will ensure that at any cut along the chain, the Hilbert space of the subsystems on each side have dimension of at least D . We can achieve this by grouping s particles at each end of the chain into a single particle, where s is chosen to be the smallest integer such that $d^s \geq D$. Denote d^s as d_{end} , the dimensionality of each of those end particles. Note that $d_{\text{end}} = d^s \leq Dd$. The dimension of the rest of the particles will remain d . We renumber the particles after the grouping, so that the new $H_{1,2}$ is now the sum of the old $H_{i,i+1}$ for i ranging from 1 to s . The term in the Hamiltonian for the last two particles is adjusted in a similar manner. We will assume from now on that the Hamiltonian is given in this form.

Let us now see how the canonical conditions can be stated in a local manner. Graphically, the second condition is equivalent to

$$\langle \mathbf{L}_{\beta}^{[j]} | \mathbf{L}_{\alpha}^{[j]} \rangle = \begin{array}{c} \text{r}^{[1]} \quad \lambda^{[2]} \quad \dots \quad \lambda^{[j-1]} \quad \Gamma^{[j-1]} \\ \circ \quad \circ \quad \dots \quad \circ \quad \circ \quad \alpha \\ \text{---} \quad \text{---} \quad \dots \quad \text{---} \quad \text{---} \\ \circ \quad \circ \quad \dots \quad \circ \quad \circ \quad \beta \end{array} = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \delta_{\alpha\beta}, \quad (2)$$

and similarly from the other side. Here the upper part of the network corresponds to $|\mathbf{L}_{\alpha}^{[j]}\rangle$, and the lower part corresponds to $|\mathbf{L}_{\beta}^{[j]}\rangle$. Notice that the canonical conditions imply that we can ‘collapse’ the network both from the left side and from the right side. Moreover, as this condition holds at every bond, it is not difficult to see that a necessary and sufficient condition for an MPS to be canonical consists of the following local conditions on $(\lambda^{[j]}, \Gamma^{[j]}, \lambda^{[j+1]})$: For every $j = 2, \dots, n - 1$,

$$\langle (\lambda^{[j]} \Gamma^{[j]})_{\alpha} | (\lambda^{[j]} \Gamma^{[j]})_{\beta} \rangle = \delta_{\alpha\beta} \quad (\text{left canonical}), \quad (3)$$

$$\langle (\Gamma^{[j]} \lambda^{[j+1]})_{\alpha} | (\Gamma^{[j]} \lambda^{[j+1]})_{\beta} \rangle = \delta_{\alpha\beta} \quad (\text{right canonical}). \quad (4)$$

For $j = 1$ and $j = n$, for $1 \leq \alpha, \beta \leq D$,

$$\langle \Gamma_{\alpha}^{[1]} | \Gamma_{\beta}^{[1]} \rangle = \langle \Gamma_{\alpha}^{[n]} | \Gamma_{\beta}^{[n]} \rangle = \delta_{\alpha\beta} \quad (\text{boundary canonical conditions}). \quad (5)$$

We also require that the λ 's are normalized, namely, that for every j from 2 to n ,

$$\langle \lambda^{[j]} | \lambda^{[j]} \rangle = 1. \quad (6)$$

Graphically, these conditions are summarized in Fig. 3.

Any triplet $(\lambda^{[j]}, \Gamma^{[j]}, \lambda^{[j+1]}) = (\lambda, \Gamma, \mu)$ that satisfies the normalization and the left and right canonical conditions [Eqs. (3), (4), and (6)] is called a *canonical triplet*. Such a triplet can be associated with a quantum state on three particles $|\psi\rangle = |\lambda \Gamma \mu\rangle = \sum_{\alpha, i, \beta} \lambda_{\alpha} \Gamma_{\alpha\beta}^i \mu_{\beta} |\alpha\rangle |i\rangle |\beta\rangle$, with the following properties: $\|\psi\| = 1$; the Schmidt basis of the first particle is the standard basis, with Schmidt coefficients $\{\lambda_{\alpha}\}$; and the Schmidt basis of the third particle is the standard basis, with Schmidt coefficients $\{\mu_{\beta}\}$. A canonical MPS can thus be described as a set of canonical triplets (or equivalently three-particle states) such that the right μ tensor of one state is equal to the left λ tensor of the next canonical triplet.

Instead of describing a canonical MPS in terms of canonical triplets (λ, Γ, μ) , we will often describe it using *canonical pairs* (λ, \mathbf{B}) , where

$$\mathbf{B} \stackrel{\text{def}}{=} \Gamma \mu.$$

The advantage is that for canonical MPSs, the elements in \mathbf{B} are always bounded (since the L_2 norm of \mathbf{B} satisfies $\|\mathbf{B}\| = \sqrt{D}$; see Sec. IID), unlike Γ whose entries can approach infinity when the corresponding μ entries approach zero.

An MPS that is described by the contraction $\Gamma^{[1]} \lambda^{[2]} \Gamma^{[2]} \lambda^{[3]} \dots \lambda^{[n]} \Gamma^{[n]}$ can also be denoted as $\Gamma^{[1]} \lambda^{[2]} \mathbf{B}^{[2]} \mathbf{B}^{[3]} \dots \mathbf{B}^{[n-1]} \Gamma^{[n]}$. No information is lost since μ can always be recovered from (λ, \mathbf{B}) : μ_{β} is the norm (see Sec. IID) of the tensor state $(\lambda \mathbf{B})_{\beta}$.¹

$$\mu_{\beta} = \left(\sum_{i, \alpha} |\lambda_{\alpha} B_{\alpha\beta}^i|^2 \right)^{1/2}.$$

We define $\mu \stackrel{\text{def}}{=} \mu(\lambda, \mathbf{B})$ this way also for noncanonical pairs.

¹Recall that μ_{β} corresponds to a Schmidt coefficient in a Schmidt decomposition that coincides with the standard basis.





FIG. 3. (a) The normalization condition for $j = 2, \dots, n$. (b) The left-right canonical conditions for $j = 2, \dots, n - 1$ [see Eqs. (3) and (4)]. (c) The boundary canonical conditions for $j = 1$ and $j = n$ [see Eq. (5)].

The advantage of working with the canonical form is that the energy of local Hamiltonians involves only the local tensors, as the following figure illustrates:



The preceding equality was obtained using the canonical conditions that are described in Eq. (2). Consequently, the energy $\langle \psi | H_{j-1,j} | \psi \rangle$ only involves five tensors: $\lambda^{[j-1]}$, $\Gamma^{[j-1]}$, $\lambda^{[j]}$, $\Gamma^{[j]}$, and $\lambda^{[j+1]}$. Similarly, $H_{1,2}$ only depends on $\Gamma^{[1]}$, $\lambda^{[2]}$, $\Gamma^{[2]}$, $\lambda^{[3]}$, and $H_{n-1,n}$ only depends on $\lambda^{[n-1]}$, $\Gamma^{[n-1]}$, $\lambda^{[n]}$, and $\Gamma^{[n]}$. It is important that each energy term does not involve tensors further to the right in the chain since the algorithm attempts to compute (or approximate) the optimal MPS up to a certain point. We would like to be able to grow the description of the state from left to right, without affecting the energies we have already computed. If matrices in the right side of the chain affected energies of terms in the left side, we would need to go back and change the MPS matrices of the particles we have already handled after we make new assignments to particles on the right. This would ruin the entire idea of dynamic programming.

Fortunately, any MPS representing a normalized state can be written as a canonical MPS with no increase in bond dimension. This follows from Ref. [20], in which it is shown that any state with Schmidt rank of at most D across any cut can be written as a canonical MPS with bond dimension D .

D. Tensor norms and distances

We use the L_2 norm on tensors $\|X\|^2 \stackrel{\text{def}}{=} \sum_{i_1, \dots, i_k} |X_{i_1 \dots i_k}|^2$. This norm of course induces a metric, namely, a way of defining the distance between tensors of the same rank. It is easy to see that the norm of a tensor C is equal to the Euclidian norm of its corresponding state $|C\rangle$. Also, for a rank-2 tensor (which can be viewed as a matrix), it is known that its operator norm is not larger than its tensor norm (which in this case is simply the Frobenius norm).

It is true that for any three tensors, B_1, B_2, B , we have $\|B_1 B - B_2 B\| \leq \|B_1 - B_2\| \cdot \|B\|$. In fact, many times in the context of MPSs, a much stronger inequality holds. Assume B connects with B_1 or B_2 along one edge, indexed by α . Assume further that $\|B_\alpha\| = 1$ for every α (in the context of canonical MPSs, it will often be the case that we consider the contraction of one side of the chain with a fixed index α of the cut edge, and this contraction is indeed of norm 1 by the canonical conditions). In this case, we have a much stronger inequality, which can be easily verified:

$$\|B_1 B - B_2 B\| = \|B_1 - B_2\|. \quad (7)$$

We can apply this to cases of interest when we compare contractions of tensors which differ in only a single term. For example, consider vector λ_α with norm 1 and two tensors $B_{\alpha, i_1, \dots, i_k}$ and $A_{\alpha, j_1, \dots, j_l}$ such that when α is fixed, the resulting tensors A_α and B_α have norm 1. Let $\hat{\lambda}$, \hat{A} , and \hat{B} be tensors with the same rank and dimensions as λ , A , and B . We have, by Eq. (7),

$$\|A \lambda B - A \hat{\lambda} B\| = \|\lambda - \hat{\lambda}\|, \quad (8)$$

and also

$$\|A \lambda B - \hat{A} \lambda B\| = \|A \lambda - \hat{A} \lambda\|. \quad (9)$$

And similarly,

$$\begin{aligned} \|A \lambda B - A \lambda \hat{B}\| &= \|\lambda B - \lambda \hat{B}\| \\ &= \left[\sum_{\alpha} |\lambda_{\alpha}|^2 \|B_{\alpha} - \hat{B}_{\alpha}\|^2 \right]^{1/2} \\ &\leq \max_{\alpha} \|B_{\alpha} - \hat{B}_{\alpha}\|. \end{aligned} \quad (10)$$

III. THE ALGORITHM

As discussed earlier, in order to carry out the outline described in Sec. IA, we would like to work with canonical MPSs. Additionally, since the tensor pairs (λ, B) for neighboring nodes overlap, we would like an ϵ net over canonical pairs such that $\mu(\lambda, B)$ could be equal to the λ of the next pair (we

call this *perfect stitching*). We do not know how to efficiently construct an ϵ net that satisfies those conditions exactly; we resort to approximately canonical MPSs with approximate stitching.

A. ϵ nets

We fix $\epsilon > 0$ (to be determined later) and define two ϵ nets. We start with discretizing $\Gamma^{[1]}$ and $\Gamma^{[n]}$.

Definition 1 (the $G_{\text{end}} \in \text{net}$). G_{end} is a set of canonical boundary tensors [see Eq. (5)] such that for any canonical boundary tensor $\hat{\Gamma}$, there is $\Gamma \in G_{\text{end}}$ such that for each α , $\|\hat{\Gamma}_\alpha - \Gamma_\alpha\| \leq \epsilon$.

We now define an ϵ net over the intermediate tensors or, more precisely, for the pairs (λ, \mathbf{B}) .

Definition 2 (the $G \in \text{net}$). G is a set of pairs of tensors (λ, \mathbf{B}) such that:

- (1) λ is positive and normalized. For all $\alpha \lambda_\alpha \geq 0$ and $\langle \lambda | \lambda \rangle = 1$.
- (2) G is an ϵ net. For every canonical triplet $(\hat{\lambda}, \hat{\Gamma}, \hat{\mu})$, there is $(\lambda, \mathbf{B}) \in G$ such $\|\hat{\lambda} \hat{\Gamma} \hat{\mu} - \lambda \mathbf{B}\| \leq \epsilon$.
- (3) \mathbf{B} is perfectly right canonical. For every α, α' , $\langle \mathbf{B}_\alpha | \mathbf{B}_{\alpha'} \rangle = \delta_{\alpha\alpha'}$ (here α, α' are the left Greek indices of \mathbf{B}).
- (4) (λ, \mathbf{B}) are approximately left canonical. For every $\beta \neq \beta'$,

$$|\langle (\lambda \mathbf{B})_\beta | (\lambda \mathbf{B})_{\beta'} \rangle| \leq 3\epsilon. \quad (11)$$

B. ϵ -net generators

We now explain how to construct such nets efficiently. Both generators for the ϵ nets will make use of the following general lemma.

Lemma 1. For any positive integers $a \leq b$ and any v in the range $(0, 1/\sqrt{a}]$, we can generate a set of $a \times b$ matrices S_{ab} over the complex numbers such that for any $A \in S_{ab}$, the rows of A are an orthonormal set of length b vectors. Furthermore, for any $a \times b$ matrix B whose rows form a set of orthonormal vectors, there is a matrix $A \in S_{ab}$ such that each row of $A - B$ has L_2 norm at most v . The size of S_{ab} is at most $O((72b/v)^{2ab})$. The time to generate S_{ab} is $O(a^2 b (72b/v)^{2ab})$. If $a = 1$, we can generate a set of vectors with real nonnegative entries, rather than complex. The size of the net is $O((72b/v)^b)$ and the time to generate it is $O(b(72b/v)^b)$.

The proof appears in the Appendix.

1. Generating G_{end}

Invoke Lemma 1 with $v = \epsilon$, $a = D$, and $b = d_{\text{end}}$. For every $A \in S_{D, d_{\text{end}}}$, add a Γ to the ϵ net, where $A_{\alpha, i} = \Gamma_\alpha^i$. Note that the conditions of Lemma 1 are satisfied if $\epsilon \leq 1/\sqrt{D}$. Since $d_{\text{end}} \leq Dd$, the size of the net is at most $(72Dd/\epsilon)^{2dD^2}$ and the time to generate it is $O(dD^3)$ times the size of the set.

2. Generating G

We generate G by first generating an $\epsilon/2$ net over the λ 's and an $\epsilon/2$ net over the \mathbf{B} 's. To generate the net of the λ 's, invoke Lemma 1 with $a = 1$, $b = D$, and the v in the lemma set to $\epsilon/2$. Note that we would like to have a λ with nonnegative real entries. According to Lemma 1, this actually requires fewer

items in our net since we are omitting the phases in each entry in the tensor. The resulting net for the λ 's has size $(144D/\epsilon)^D$ and can be generated in time $O(D(144D/\epsilon)^D)$.

To generate the net over the \mathbf{B} 's, we invoke Lemma 1 with $a = D$, $b = dD$, and $v = \epsilon/2$. Note that in order to invoke Lemma 1, we require that $\epsilon \leq 2/\sqrt{D}$. For any matrix $A_{\alpha, (i, \beta)}$ in the set, we generate a tensor \mathbf{B} where $B_{\alpha, \beta}^i = A_{\alpha, (i, \beta)}$. This way we generate a set of pairs (λ, \mathbf{B}) which satisfies both the normalization condition [condition (1) of Definition 2] and the condition of being perfect right canonical [condition (3) of Definition 2].

To see that we in fact have an ϵ net [i.e., condition (2) is satisfied], consider a perfectly canonical pair (λ, \mathbf{B}) , and let us find a pair $(\hat{\lambda}, \hat{\mathbf{B}})$ in the net that is ϵ close to it. We first replace λ with a $\hat{\lambda}$ from the first net and then replace \mathbf{B} with a $\hat{\mathbf{B}}$ from the second net. Using Eq. (8), we have that

$$\|\lambda \mathbf{B} - \hat{\lambda} \hat{\mathbf{B}}\| = \|\lambda - \hat{\lambda}\| \leq \frac{\epsilon}{2}.$$

Using Eq. (10), we also have

$$\|\hat{\lambda} \hat{\mathbf{B}} - \hat{\lambda} \mathbf{B}\| \leq \max_\alpha \|\mathbf{B}_\alpha - \hat{\mathbf{B}}_\alpha\| \leq \frac{\epsilon}{2}.$$

Next, we discard all tensors (λ, \mathbf{B}) that are not approximately left canonical, namely, those that violate condition (4). It remains to show that the remaining tensors still satisfy condition (2), that is, the ϵ -net condition. We do that by showing that a pair (λ, \mathbf{B}) that is ϵ close to a canonical triplet must necessarily be approximately left canonical. Therefore, such a pair would not have been eliminated.

To see this, let the tensor $\mathbf{A} = \lambda \Gamma \mu$ be the contraction of the canonical triplet and \mathbf{C} be the contraction of $\hat{\lambda} \hat{\mathbf{B}}$ from the net such that $\|\mathbf{A} - \mathbf{C}\| \leq \epsilon$. The fact that \mathbf{A} is perfectly left canonical is expressed in the fact that for every $\beta \neq \beta'$, $\langle \mathbf{A}_\beta | \mathbf{A}_{\beta'} \rangle = 0$. To prove that \mathbf{C} is approximately left canonical, we need to show $|\langle \mathbf{C}_\beta | \mathbf{C}_{\beta'} \rangle| \leq 3\epsilon$. Indeed, $\|\mathbf{A} - \mathbf{C}\| \leq \epsilon$ implies $\|\mathbf{A}_\beta - \mathbf{C}_\beta\| \leq \epsilon$ for every β . Assume $\beta \neq \beta'$. Then

$$\begin{aligned} |\langle \mathbf{C}_\beta | \mathbf{C}_{\beta'} \rangle| &= |\langle \mathbf{A}_\beta + (\mathbf{C}_\beta - \mathbf{A}_\beta) | \mathbf{A}_{\beta'} + (\mathbf{C}_{\beta'} - \mathbf{A}_{\beta'}) \rangle| \\ &\leq |\langle \mathbf{A}_\beta | \mathbf{A}_{\beta'} \rangle| + |\langle \mathbf{A}_\beta | \mathbf{C}_{\beta'} - \mathbf{A}_{\beta'} \rangle| \\ &\quad + |\langle \mathbf{C}_\beta - \mathbf{A}_\beta | \mathbf{A}_{\beta'} \rangle| + |\langle \mathbf{C}_\beta - \mathbf{A}_\beta | \mathbf{C}_{\beta'} - \mathbf{A}_{\beta'} \rangle| \\ &\leq \|\mathbf{A}_\beta\| \|\mathbf{C}_{\beta'} - \mathbf{A}_{\beta'}\| + \|\mathbf{A}_{\beta'}\| \|\mathbf{C}_\beta - \mathbf{A}_\beta\| \\ &\quad + \|\mathbf{C}_\beta - \mathbf{A}_\beta\| \|\mathbf{C}_{\beta'} - \mathbf{A}_{\beta'}\| \\ &\leq 2\epsilon + \epsilon^2 \leq 3\epsilon. \end{aligned}$$

This concludes the proof that G is indeed an ϵ net according to Definition 2.

3. Complexity of generating G and G_{end}

By Lemma 1, $N \stackrel{\text{def}}{=} |G|$, the size of the ϵ net G is

$$N = O\left(\frac{144dD}{\epsilon}\right)^{D+2dD^2}. \quad (12)$$

This is the size of the set formed by taking all pairs (λ, \mathbf{B}) , where each λ and \mathbf{B} come from their respective nets. The time required to generate the original net (before tensors are discarded) is $O(dD^3 N)$. The cost of checking whether a (λ, \mathbf{B}) pair is approximately left canonical is $O(dD^3)$, so the total cost of generating the net is $O(dD^3 N)$.

For G_{end} , both the number of points and the running time which were determined in Sec. III B 1, are bounded above by the corresponding bounds of G .

C. The algorithm

When processing particle j , the algorithm creates a list L_j of partial solutions, one for each (λ, \mathbf{B}) pair in G . For each such partial solution, a tail (i.e., the tensors to the left of the j th particle) and energy is kept.

First step. Create the first list L_2 : For each $(\lambda^{[2]}, \mathbf{B}^{[2]}) \in G$, find its tail, namely, the $\Gamma^{[1]} \in G_{\text{end}}$ which minimizes the energy with respect to $H_{1,2}$ of the tensor $\Gamma^{[1]} \lambda^{[2]} \mathbf{B}^{[2]}$. Denote this minimal energy as $E_2(\lambda^{[2]}, \mathbf{B}^{[2]})$. We keep both the tail and the computed energy, for each pair $(\lambda^{[2]}, \mathbf{B}^{[2]}) \in G$.

Going from $j = 3$ to $j = n - 1$. We assume we have created the list L_{j-1} . For each pair $(\lambda^{[j-1]}, \mathbf{B}^{[j-1]}) \in G$ there is a tail in L_{j-1} ,

$$\Gamma^{[1]}, (\lambda^{[2]}, \mathbf{B}^{[2]}), (\lambda^{[3]}, \mathbf{B}^{[3]}), \dots, (\lambda^{[j-2]}, \mathbf{B}^{[j-2]}),$$

and an energy value that we denote as $E_{j-1}(\lambda^{[j-1]}, \mathbf{B}^{[j-1]})$. To create L_j , we find a tail for each $(\lambda^{[j]}, \mathbf{B}^{[j]}) \in G$. We require that the tail for a given $(\lambda^{[j]}, \mathbf{B}^{[j]})$ is an item in L_{j-1} which satisfies the ‘‘stitching’’ condition:

$$\|\mu(\lambda^{[j-1]}, \mathbf{B}^{[j-1]}) - \lambda^{[j]}\| \leq 2\epsilon. \quad (13)$$

We pick the tail for $(\lambda^{[j]}, \mathbf{B}^{[j]})$ to be an item in L_{j-1} which satisfies the stitching condition and minimizes $H_{j-1,j}(\lambda^{[j-1]} \mathbf{B}^{[j-1]} \mathbf{B}^{[j]}) + E_{j-1}(\lambda^{[j-1]}, \mathbf{B}^{[j-1]})$. The minimum such value is defined to be $E_j(\lambda^{[j]}, \mathbf{B}^{[j]})$.

Final step. The final step, $j = n$, is exactly as in the intermediate steps except the algorithm goes over $\Gamma^{[n]} \in G_{\text{end}}$, rather than over pairs from G and there is no stitching constraint. More precisely, we pick the tail for $\Gamma^{[n]}$ to be the item in L_{n-1} which minimizes $H_{n-1,n}(\lambda^{[n-1]} \mathbf{B}^{[n-1]} \Gamma^{[n]}) + E_{n-1}(\lambda^{[n-1]}, \mathbf{B}^{[n-1]})$. The minimal value is defined to be $E_n(\Gamma^{[n]})$.

Finally, we choose $\Gamma^{[n]}$ which minimizes $E_n(\Gamma^{[n]})$. We output the MPS that is defined by $\Gamma^{[n]}$ and its tail,

$$|\Omega\rangle \stackrel{\text{def}}{=} |\Gamma^{[1]} \lambda^{[2]} \mathbf{B}^{[2]} \mathbf{B}^{[3]} \dots \mathbf{B}^{[n-1]} \Gamma^{[n]}\rangle, \quad (14)$$

together with the energy which the algorithm calculated:

$$E_{\text{alg}}(\Omega) \stackrel{\text{def}}{=} E_n(\Gamma^{[n]}). \quad (15)$$

Note that since each $(\lambda^{[j]}, \mathbf{B}^{[j]})$ is perfectly right canonical, the state $|\Omega\rangle$ is normalized. This can be seen by contracting the tensor network corresponding to the inner product $\langle \Omega | \Omega \rangle$ from right to left.

Unlike in the classical case, our algorithm does not search all states due to the discretization. Moreover, it does not optimize over the real energy of the states that it does check, but rather over $E_{\text{alg}}(\Omega) = \sum_j H_{j-1,j}(\lambda^{[j-1]} \mathbf{B}^{[j-1]} \mathbf{B}^{[j]})$. E_{alg} is different from the true energy E because the states are not exactly canonical. Note that the output $E_{\text{alg}}(\Omega)$ is thus just an approximation of the real energy $E(\Omega)$ of the output MPS $|\Omega\rangle$. We output $E_{\text{alg}}(\Omega)$ anyway, since our guarantee on its error is somewhat better than on the error for $E(\Omega)$, as we see in Sec. IV.

The following claim easily follows from the same reasoning as for the classical dynamic programming algorithm:

Claim 1. The algorithm finds the state which minimizes E_{alg} among all MPSs of the form $\Gamma^{[1]} \lambda^{[2]} \mathbf{B}^{[2]} \mathbf{B}^{[3]} \dots \mathbf{B}^{[n-1]} \Gamma^{[n]}$ such that $\Gamma^{[1]}, \Gamma^{[n]} \in G_{\text{end}}$, $(\lambda^{[j]}, \mathbf{B}^{[j]}) \in G$ for all $j \in \{2, \dots, n-1\}$, and the stitching conditions [Eq. (13)] are all satisfied.

IV. ERROR AND COMPLEXITY ANALYSIS

In order to finish the proof of Theorem 1, we will prove the theorem in what follows. As noted earlier, this theorem actually gives a better error bound on $E_{\text{alg}}(\Omega)$ than the bound on $E(\Omega)$ that is given in Theorem 1.

Theorem 3 (error bound). Let E_0 be the minimal energy that can be achieved by a state with bond dimension D and J the maximal operator norm $\|H_{j,j+1}\|$ over all terms. Then

$$E_{\text{alg}}(\Omega) - 6Jn\epsilon \leq E_0 \leq E(\Omega) \leq E_{\text{alg}}(\Omega) + \frac{3}{2}JD^2n^2\epsilon. \quad (16)$$

It is easy to verify that as long as $nD^2 \geq 12$, Eq. (16) implies Eq. (1) of Theorem 1.

Proof. By definition, $E_0 \leq E(\Omega)$. We first prove that $E_{\text{alg}}(\Omega) - 6Jn\epsilon \leq E_0$. Let

$$|\psi\rangle = |\hat{\Gamma}^{[1]} \hat{\lambda}^{[2]} \hat{\Gamma}^{[2]} \dots \hat{\lambda}^{[n]} \hat{\Gamma}^{[n]}\rangle$$

be a state with $E(\psi) = E_0$ of bond dimension D , written as a canonical MPS. For every triplet $(\hat{\lambda}^{[j]}, \hat{\Gamma}^{[j]}, \hat{\lambda}^{[j+1]})$ for $j = 2, \dots, n-1$, we associate a pair $(\tilde{\lambda}^{[j]}, \tilde{\mathbf{B}}^{[j]}) \in G$ which is ϵ close to that triplet. In addition, we find $\tilde{\Gamma}_1 \in G_{\text{end}}$ close to $\hat{\Gamma}_1$ and $\tilde{\Gamma}_n \in G_{\text{end}}$ close to $\hat{\Gamma}_n$. We define the state

$$|\phi\rangle = |\tilde{\Gamma}_1 \tilde{\lambda}_2 \tilde{\mathbf{B}}_2 \tilde{\mathbf{B}}_3 \dots \tilde{\mathbf{B}}_{n-1} \tilde{\Gamma}_n\rangle.$$

Just like $|\Omega\rangle$, this state is normalized due to the fact that the tensors in G_{end} and G are perfectly right canonical.

To show that $E_{\text{alg}}(\Omega) - 6Jn\epsilon \leq E_0$, we will first show that $E_{\text{alg}}(\Omega) \leq E_{\text{alg}}(\phi)$ and then that $|E_{\text{alg}}(\phi) - E_0| \leq 6Jn\epsilon$. The first inequality follows from the fact that $|\phi\rangle$ belongs to the set of states over which the dynamic algorithm searches (see Claim 1), since the $\tilde{\lambda}^{[j-1]} \tilde{\mathbf{B}}^{[j-1]}$ and $\tilde{\lambda}^{[j]} \tilde{\mathbf{B}}^{[j]}$ satisfy the stitching condition (13), as promised by the following lemma.

Lemma 2. For every $j = 3, \dots, n-1$,

$$\|\mu(\tilde{\lambda}^{[j-1]}, \tilde{\mathbf{B}}^{[j-1]}) - \tilde{\lambda}^{[j]}\| \leq 2\epsilon. \quad (17)$$

Proof. We use the fact (established in Lemma 4 in the Appendix) that for any two bipartite states, $|A\rangle = \sum_i a_i |i\rangle |A_i\rangle$ with normalized $|A_i\rangle$, $|B\rangle = \sum_i b_i |i\rangle |B_i\rangle$ with normalized $|B_i\rangle$, we have $\sum_i |a_i - b_i|^2 \leq \|A - B\|^2$.

The tensors $\hat{\lambda}^{[j]} \hat{\Gamma}^{[j]} \hat{\lambda}^{[j+1]}$ and $\tilde{\lambda}^{[j]} \tilde{\mathbf{B}}^{[j]}$ represent two quantum states on three particles, where in both states, the Schmidt basis of the *first* particle is the standard basis, and the perfect right canonical condition of Definition 2 [or, alternatively, the condition of Eq. (4)] holds. The Schmidt coefficients are given as $\{\hat{\lambda}_\alpha^{[j]}\}$ and $\{\tilde{\lambda}_\alpha^{[j]}\}$, respectively. According to the preceding fact (Lemma 4),

$$\|\hat{\lambda}^{[j]} - \tilde{\lambda}^{[j]}\| \leq \|\hat{\lambda}^{[j]} \hat{\Gamma}^{[j]} \hat{\lambda}^{[j+1]} - \tilde{\lambda}^{[j]} \tilde{\mathbf{B}}^{[j]}\| \leq \epsilon. \quad (18)$$

Similarly, we know that $\|\hat{\lambda}^{[j-1]} \hat{\Gamma}^{[j-1]} \hat{\lambda}^{[j]} - \tilde{\lambda}^{[j-1]} \tilde{\mathbf{B}}^{[j-1]}\| \leq \epsilon$. Consider now these three-particle states expanded in terms of the basis vectors $|\beta\rangle$ of the *third* particle. Denote these expansions as $\sum_{\beta} a_{\beta} |v_{\beta}\rangle |\beta\rangle$ with normalized $|v_{\beta}\rangle$ and $\sum_{\beta} b_{\beta} |w_{\beta}\rangle |\beta\rangle$ with normalized $|w_{\beta}\rangle$, respectively. Then, by definition, $a_{\beta} = \hat{\lambda}_{\beta}^{[j]}$ and $b_{\beta} = \mu_{\beta}(\hat{\lambda}^{[j-1]}, \tilde{\mathbf{B}}^{[j-1]})$. We can therefore apply again Lemma 4 and get $\|\mu(\hat{\lambda}^{[j-1]}, \tilde{\mathbf{B}}^{[j-1]}) - \hat{\lambda}^{[j]}\| \leq \epsilon$. Together with Eq. (18), we therefore obtain $\|\mu(\hat{\lambda}^{[j-1]}, \tilde{\mathbf{B}}^{[j-1]}) - \tilde{\lambda}^{[j]}\| \leq 2\epsilon$. ■

Thus far, we have established that $E_{\text{alg}}(\Omega) \leq E_{\text{alg}}(\phi)$. To finish the proof that $E_{\text{alg}}(\Omega) - 6Jn\epsilon \leq E_0$, we will show that $|E_{\text{alg}}(\phi) - E_0| \leq 6Jn\epsilon$. Observe that each energy term in E_0 depends solely on two overlapping triplets $\hat{\lambda}^{[j]} \hat{\Gamma}^{[j]} \hat{\lambda}^{[j+1]}$ and $\hat{\Gamma}^{[j+1]} \hat{\lambda}^{[j+1]} \hat{\Gamma}^{[j+2]}$. The corresponding energy term in $E_{\text{alg}}(\phi)$ depends only on $\hat{\lambda}^{[j]} \tilde{\mathbf{B}}^{[j]} \tilde{\mathbf{B}}^{[j+1]}$. We now bound the distance between these two tensors. We have

$$\begin{aligned} & \|\hat{\lambda}^{[j]} \tilde{\mathbf{B}}^{[j]} \tilde{\mathbf{B}}^{[j+1]} - \hat{\lambda}^{[j]} \hat{\Gamma}^{[j]} \hat{\lambda}^{[j+1]} \hat{\Gamma}^{[j+1]} \hat{\lambda}^{[j+2]}\| \\ &= (\hat{\lambda}^{[j]} \tilde{\mathbf{B}}^{[j]} - \hat{\lambda}^{[j]} \hat{\Gamma}^{[j]} \hat{\lambda}^{[j+1]}) \tilde{\mathbf{B}}^{[j+1]} \\ & \quad + \hat{\lambda}^{[j]} \hat{\Gamma}^{[j]} (\hat{\lambda}^{[j+1]} - \tilde{\lambda}^{[j+1]}) \tilde{\mathbf{B}}^{[j+1]} \\ & \quad + \hat{\lambda}^{[j]} \hat{\Gamma}^{[j]} (\tilde{\lambda}^{[j+1]} \tilde{\mathbf{B}}^{[j+1]} - \hat{\lambda}^{[j+1]} \hat{\Gamma}^{[j+1]} \hat{\lambda}^{[j+2]}). \end{aligned}$$

Taking the left-hand side and right-hand side of the preceding equation, and using Eqs. (8) and (9), we have that

$$\begin{aligned} & \|\hat{\lambda}^{[j]} \tilde{\mathbf{B}}^{[j]} \tilde{\mathbf{B}}^{[j+1]} - \hat{\lambda}^{[j]} \hat{\Gamma}^{[j]} \hat{\lambda}^{[j+1]} \hat{\Gamma}^{[j+1]} \hat{\lambda}^{[j+2]}\| \\ & \leq \|\hat{\lambda}^{[j]} \tilde{\mathbf{B}}^{[j]} - \hat{\lambda}^{[j]} \hat{\Gamma}^{[j]} \hat{\lambda}^{[j+1]}\| + \|\hat{\lambda}^{[j+1]} - \tilde{\lambda}^{[j+1]}\| \\ & \quad + \|\hat{\lambda}^{[j+1]} \tilde{\mathbf{B}}^{[j+1]} - \hat{\lambda}^{[j+1]} \hat{\Gamma}^{[j+1]} \hat{\lambda}^{[j+2]}\|. \end{aligned}$$

The first and third terms in the above sum can be bounded by ϵ because of the condition of the ϵ net G . The norm of the middle term is bounded in Eq. (18). Therefore, the norm of the difference between the tensors is at most 3ϵ . It follows that the difference between the two energy contributions is at most $6\epsilon \|H_{j,j+1}\| \leq 6\epsilon J$.


We illustrate the boundary cases by working through the analysis for the left end of the chain. We want to bound $\|\hat{\Gamma}^{[1]} \hat{\lambda}^{[2]} \hat{\Gamma}^{[2]} \hat{\lambda}^{[3]} - \tilde{\Gamma}^{[1]} \tilde{\lambda}^{[2]} \tilde{\mathbf{B}}^{[2]}\|$. Note that $\|\hat{\Gamma}^{[1]} (\hat{\lambda}^{[2]} \hat{\Gamma}^{[2]} \hat{\lambda}^{[3]} - \tilde{\lambda}^{[2]} \tilde{\mathbf{B}}^{[2]})\|$ is bounded by ϵ because of the conditions on the ϵ net and Eq. (7). Using Eq. (10), we have that

$$\|(\hat{\Gamma}^{[1]} - \tilde{\Gamma}^{[1]}) \tilde{\lambda}^{[2]} \tilde{\mathbf{B}}^{[2]}\| \leq \max_{\alpha} \|\hat{\Gamma}_{\alpha}^{[1]} - \tilde{\Gamma}_{\alpha}^{[1]}\| \leq \epsilon.$$


Hence, the overall bound on the difference is 2ϵ . It follows that the difference between the two energy contributions is at most $4\epsilon \|H_{1,2}\| \leq 4\epsilon J$. A similar argument holds for $H_{n-1,n}$.

We now prove the right inequality in Theorem 3 by showing that $|E(\Omega) - E_{\text{alg}}(\Omega)| \leq \frac{3}{2} J D^2 n^2 \epsilon$. We bound the difference in energy for each term $H_{j-1,j}$. The contribution of this term to $E_{\text{alg}}(\Omega)$ is calculated from $\lambda^{[j-1]} \mathbf{B}^{[j-1]} \mathbf{B}^{[j]}$. The true energy, however, depends on $\Gamma^{[1]} \lambda^{[2]} \mathbf{B}^{[2]} \mathbf{B}^{[3]} \dots \mathbf{B}^{[j]}$ since $|\Omega\rangle$ is only approximately left canonical. We will show that the error accumulates linearly as we sweep from left to right, summing up to $3j J D^2 \epsilon$ for $H_{j-1,j}$. Therefore, the total error is $|E_{\text{alg}}(\Omega) - E(\Omega)| \leq \frac{3}{2} J D^2 n^2 \epsilon$.

We now provide a more accurate argument. The energy estimate for the term $H_{j-1,j}$ is calculated from the contraction $\lambda^{[j-1]} \mathbf{B}^{[j-1]} \mathbf{B}^{[j]}$. Graphically, this contribution is given as



The true energy, however, is calculated from the contraction of $\lambda^{[2]} \mathbf{B}^{[2]} \mathbf{B}^{[3]} \dots \mathbf{B}^{[j]}$. Graphically, this is given as



[Notice that we have collapsed the $\Gamma^{[1]}$ terms because of the canonical condition (5); see Fig. 3(b).]

Had the state $|\Omega\rangle$ been perfectly left canonical, the two would have been the same. However, since it is only approximately canonical from the left, there is some difference that can be bounded. The analysis is done iteratively from left to right. We start by writing

$$\lambda^{[2]} \begin{array}{c} \beta \\ \beta' \end{array} = \mu^{[2]} \begin{array}{c} \beta \\ \beta' \end{array} + R \begin{array}{c} \beta \\ \beta' \end{array} = \lambda^{[3]} \begin{array}{c} \beta \\ \beta' \end{array} + \Delta \begin{array}{c} \beta \\ \beta' \end{array}$$


In this picture, the tensor $R_{\beta\beta'}$ is defined to be off diagonal (i.e., equal to zero on the diagonal: $R_{\beta\beta} = 0$) and for the $\beta \neq \beta'$ terms, it is defined by $R_{\beta\beta'} = \langle (\lambda^{[2]} \mathbf{B}^{[2]})_{\beta} | (\lambda^{[2]} \mathbf{B}^{[2]})_{\beta'} \rangle = \sum_{\alpha,i} |\lambda_{\alpha}^{[2]}|^2 B_{\alpha\beta}^{[2]i} (B_{\alpha\beta'}^{[2]i})^*$. Δ is defined as

$$\Delta_{\beta\beta'} \stackrel{\text{def}}{=} R_{\beta\beta'} + \delta_{\beta\beta'} (|\lambda_{\beta}^{[3]}|^2 - |\mu_{\beta}^{[2]}|^2).$$


Using the fact that $(\lambda^{[2]}, \mathbf{B}^{[2]})$ is approximately left canonical [see Eq. (11)] and the stitching conditions of $\lambda^{[3]}$ and $\mu^{[2]}$ [see Eq. (13)], it is easy to see that for every β, β' ,

$$|\Delta_{\beta\beta'}| \leq 3\epsilon. \quad (19)$$

We may therefore write the true energy contribution as the sum of



and



The analysis of the first term is done in the next iteration step. The second term can be seen as the error introduced by the fact that $(\lambda^{[2]}, \mathbf{B}^{[2]})$ is approximately left canonical. To estimate its size, notice that it can be viewed as the expectation value of

the operator $\mathbf{\Delta} \otimes H_{j-1,j}$ (here $\mathbf{\Delta}$ is viewed as a matrix), with respect to the MPS that is described by $|\mathbf{B}^{[3]}\mathbf{B}^{[4]}\dots\mathbf{B}^{[j]}\rangle$. Using Eq. (19) and the assumption $\|H_{j-1,j}\| \leq J$, it is easy to see that $\|\mathbf{\Delta} \otimes H_{j-1,j}\| \leq 3JD\epsilon$. Here, in both cases, we used $\|\cdot\|$ to denote the *operator* norm of $\mathbf{\Delta} \otimes H_{j-1,j}$, instead of the usual tensor norm; we can do this since the operator norm is at most as large as the tensor norm, and the tensor norm of $\mathbf{\Delta}$ is at most $3D\epsilon$. Moreover, the norm of the MPS $|\mathbf{B}^{[3]}\mathbf{B}^{[4]}\dots\mathbf{B}^{[j]}\rangle$ is exactly \sqrt{D} (it would have been exactly 1 had there been a $\lambda^{[3]}$ term before $\mathbf{B}^{[3]}$), and therefore the amplitude of second term is upper bounded by $3JD^2\epsilon$.

Carrying the same analysis all way to $(\lambda^{[j-2]}, \mathbf{B}^{[j-2]})$, we end up with a term that is identical to the energy estimation of the algorithm, plus some error term whose amplitude is at most $3jJD^2\epsilon$. Therefore, by simple algebra, we have that for the total system,

$$|E_{\text{alg}} - E(\Omega)| \leq \frac{3}{2}JD^2n^2\epsilon. \quad (20)$$

This concludes the proof of Theorem 3. \blacksquare

For a target error δ , we select $\epsilon \leq \frac{\delta}{2JD^2n^2}$. Using the bound from Eq. (12), we get that the size of the net for the interior particles is

$$N = O\left(\left(\frac{144JD^3n^2}{\delta}\right)^{D+2dD^2}\right). \quad (21)$$

Note that in using Lemma 1, we required that $\epsilon \leq 1/\sqrt{D}$. It is reasonable to expect that $\delta/Jn < 1$ (meaning that the desired error is at most the maximum energy in the system), which implies that this condition is met. The algorithm has n iterations in which $O(N^2)$ possible extensions for the MPS are considered. For each such possibility, we perform a contraction of tensors $(\lambda, \mathbf{B}, \mathbf{B}')$ in order to evaluate the energy of a particular term. This contraction takes time $O(D^3d^2)$. Thus, the total running time is $O(nN^2D^3d^2)$.

V. COMMUTING HAMILTONIAN IN 1D

We now prove Theorem 2. Let us first notice that Theorem 1 immediately implies the first claim in Theorem 2, namely, that approximating the ground state and ground energy of a commuting Hamiltonian in 1D to within polynomially good accuracy can be done efficiently. This follows from the well-known fact that the ground state of a commuting Hamiltonian in 1D can be described by an MPS of constant bond dimension. We can therefore apply Theorem 1 to the problem, and hence approximate both the ground state and the ground energy efficiently.

For completeness, here is a sketch of a proof of this fact: Assume we have a 2-local commuting Hamiltonian in 1D. If the Hamiltonian is k -local for $k > 2$, just combine adjacent particles together. To see that there is a ground state which is described by an MPS of constant bond dimension, notice that for any commuting Hamiltonian, there is a ground state $|\psi\rangle$ which is an eigenvector of each of the terms in the Hamiltonian, with some well defined eigenvalue for each term. For each term, consider the projection onto the eigenspace corresponding to that eigenvalue. For any state with nonzero projection on the ground state, applying these projections (no matter the order) would result in a ground state. Since there

is always a computational basis state $|w\rangle$ that has a nonzero projection on the ground state, we can express a ground state as the projection of all these local terms applied to $|w\rangle$. We first apply the projections which interact the pairs of particles (1,2), (3,4), etc.; we then apply the projections that interact the pairs of particles (2,3), (4,5), etc. This sequence of operations can be viewed as a tensor network of depth 2. We can thus represent the ground state as the contraction of a tensor network of depth 2. It can be easily seen that such a state must have a constant Schmidt rank along any cut between the left and right sides; to move to an MPS of a constant bond dimension, use Vidal's result [20].

We now give the full proof of Theorem 2.

Proof: We assume that the ground. In other words, we are promised that the ground energy is separated from the rest of the eigenvalues of the Hamiltonian by a gap $\Delta \geq 1/n^c$ for some constant c . Notice that we do not assume a unique ground state.

The first step of the proof would be to use Theorem 1 to find an MPS $|\Omega\rangle$ of constant bond dimension such that $\langle\Omega|H|\Omega\rangle \leq E_0 + \Delta/3$. From the preceding discussion, it is clear that this can be done in polynomial time. Next, we would like to project this MPS sequentially on some chosen eigenspaces of the Hamiltonians along the chain. As we are in a commuting system, this would result in a common eigenvector of all Hamiltonians, and therefore an eigenvector of H itself. If we manage to do this without increasing the energy above $E_0 + \Delta$, then by the existence of the gap, we are promised to have reached a ground state.

To do this, we rely on the following lemma.

Lemma 3. Let $H = \sum_i H_i$ be a commuting local Hamiltonian system with ground energy E_0 , and let $|\psi\rangle$ be a state such that $\langle\psi|H|\psi\rangle = E_0 + h$. Consider one term H_i in H with k eigenvalues and projections P_1, \dots, P_k into the corresponding eigenspaces. For every $j = 1, \dots, k$, let $|\psi_j\rangle$ be the *normalization* of $P_j|\psi\rangle$, and let $c_j = \langle\psi|P_j|\psi\rangle$. Then for any $n > 2$ there is always a j such that $c_j \geq \frac{1}{kn^2}$ and $\langle\psi_j|H|\psi_j\rangle \leq E_0 + (1 + \frac{1}{n})h$.

Proof. As the $\{H_i\}$ terms are commuting, it follows that

$$\begin{aligned} \langle\psi|H|\psi\rangle &= \langle\psi|P_1H P_1|\psi\rangle + \langle\psi|P_2H P_2|\psi\rangle \\ &\quad + \dots + \langle\psi|P_kH P_k|\psi\rangle \\ &= c_1\langle\psi_1|H|\psi_1\rangle + c_2\langle\psi_2|H|\psi_2\rangle \\ &\quad + \dots + c_k\langle\psi_k|H|\psi_k\rangle, \end{aligned}$$

with $\sum_{j=1}^k c_j = 1$. Assume, by contradiction, that for every j , either $c_j < \frac{1}{kn^2}$ or $\langle\psi_j|H|\psi_j\rangle > E_0 + (1 + \frac{1}{n})h$. Then partition the k eigenspaces into two subsets: subset A , in which the first condition holds, and subset B , in which the second condition holds. Then

$$\begin{aligned} E_0 + h &= \langle\psi|H|\psi\rangle \\ &= \sum_A c_j \langle\psi_j|H|\psi_j\rangle + \sum_B c_j \langle\psi_j|H|\psi_j\rangle \\ &\geq E_0 \sum_A c_j + \left[E_0 + \left(1 + \frac{1}{n}\right)h\right] \sum_B c_j \\ &= E_0 + \left(1 + \frac{1}{n}\right)h \sum_B c_j, \end{aligned}$$

using $\sum_j c_j = 1$. Since $\sum_A c_j \leq \frac{k}{kn^2} = \frac{1}{n^2}$, we have that $\sum_B c_j \geq 1 - \frac{1}{n^2}$. Plugging this into the preceding equality implies $h > h(1 + \frac{1}{n})(1 - \frac{1}{n^2})$, which is a contradiction for $n > 2$. ■

We now apply the lemma sequentially to project the approximate state $|\Omega\rangle$ on the relevant local eigenspaces. We start with $H_{1,2}$, where we use $h = \Delta/3$ in the lemma. The lemma promises the existence of a subspace indexed j (out of k possible j 's) which, if we project $|\Omega\rangle$ onto that subspace, the projection will not have too-large energy. We denote c_j and P_j as c_{12} and P_{12} , respectively. (In what follows, we explain how all calculations required for finding the promised j can be done efficiently.) We proceed to find c_{23} and P_{23} for the next term $H_{2,3}$, using the newly projected state, and so on up to $H_{n,n-1}$. After applying the $n - 1$ projections, using the lemma $n - 1$ times, we arrive at a state $|\psi\rangle$ given as

$$|\psi\rangle = \frac{1}{\sqrt{c_{12}c_{23}\cdots c_{n-1,n}}} P_{12}P_{23}\cdots P_{n-1,n}|\Omega\rangle,$$

which satisfies

$$\langle\psi|H|\psi\rangle \leq E_0 + \left(1 + \frac{1}{n}\right)^{n-1} \frac{\Delta}{3} \leq E_0 + \frac{e\Delta}{3}.$$

Using the assumption of the gap and the fact that $|\psi\rangle$ is an eigenvector of H , it must be that $|\psi\rangle$ is a ground state and $\langle\psi|H|\psi\rangle = E_0$.

We now argue why finding the j whose existence is promised by the lemma can be done efficiently. Consider, for example, the term $H_{m,m+1}$. To find the relevant j , we have to compute, for the current state $|\psi\rangle$, both the norms squared $c_j = \langle\psi|P_j|\psi\rangle$ and the expectation values $\langle\psi_j|H|\psi_j\rangle = \frac{1}{c_j}\langle\psi|P_jHP_j|\psi\rangle$ for all eigenspaces P_j of $H_{m,m+1}$. Note first that we are handling here real numbers; the projections P_j on the eigenspaces of $H_{m,m+1}$ may require infinite precision to describe exactly in binary (or any other) representation. We truncate the entries in the projections to exponentially good precision, using polynomially many bits, so that all the calculations can be performed efficiently. This introduces an exponentially small error.

The expressions we are interested in calculating are all of the form

$$\langle\Omega|P_{12}\cdots P_{m,m-1}\cdot P_jOP_j\cdot P_{m,m-1}\cdots P_{12}|\Omega\rangle, \quad (22)$$

where O can be either a local Hamiltonian $H_{i,i+1}$ or the identity and the $P_{i,i+1}$ are projections on eigenspaces of the local terms. Recalling that $|\Omega\rangle$ is a constant-bond MPS, and using the fact that the projections commute between themselves, we can write Eq. (22) as a constant depth-tensor network. This is done by partitioning the projections into two layers: In one layer the projections that work on the sites $(1,2),(3,4),(5,6),\dots$, and on the other, the projections that act on the sites $(2,3),(4,5),(6,7),\dots$. The resultant tensor network is shown in Fig. 4. One dimensional tensor networks with constant depth can be efficiently calculated on a classical computer because




FIG. 4. An illustration of how the expression in Eq. (22) is given as a tensor network with a constant number of horizontal layers. Specifically, the figure describes the tensor network of $\langle\Omega|P_{12}P_{23}\cdots H_{j,j+1}\cdots P_{23}P_{12}|\Omega\rangle$.

their bubble width is constant when swallowed from left to right [18].²

Thus, all calculations (under our assumptions of polynomially many bits of precision of the P_j 's) can be performed efficiently. The resulting state is given by a tensor network of constant depth (namely the original $|\Omega\rangle$ on which the chosen projections are applied.) As before, this can be modified to a MPS of constant bond dimension using Vidal's result [20], concluding the proof. ■

A. A proof for the commuting Hamiltonians case, based on Ref. [8]

First we describe the alternate algorithm assuming we have the ability to perform arithmetic operations with infinite precision and then discuss the consequences of limited precision. Reference [8] proves certain properties about the ground states of 2-local commuting Hamiltonians in which the interaction graph is a general graph. We express those properties for the special case of interest here in which the graph is a line. Let \mathcal{H}_j be the Hilbert space of particle j . It is shown in Ref. [8] that when the terms of the Hamiltonian commute, the Hilbert space of each particle can be expressed as a direct sum, $\mathcal{H}_j = \oplus_{\alpha_j} \mathcal{H}_j^{(\alpha_j)}$, such that each $\mathcal{H}_j^{(\alpha_j)}$ can then be expressed as a tensor product of three spaces,

$$\mathcal{H}_j^{(\alpha_j)} = \mathcal{H}_{L,j}^{(\alpha_j)} \otimes \mathcal{H}_{C,j}^{(\alpha_j)} \otimes \mathcal{H}_{R,j}^{(\alpha_j)}.$$

This structure has the property that $H_{j,j+1}$ leaves the subspaces $\mathcal{H}_j^{(\alpha_j)} \otimes \mathcal{H}_{j+1}^{(\alpha_{j+1})}$ invariant, and moreover, when restricted to such a subspace, $H_{j,j+1}$ acts nontrivially only on $\mathcal{H}_{R,j}^{(\alpha_j)} \otimes \mathcal{H}_{L,j+1}^{(\alpha_{j+1})}$ [the right part of $\mathcal{H}_j^{(\alpha_j)}$ and the left part of $\mathcal{H}_{j+1}^{(\alpha_{j+1})}$]. Consequently, there exists a ground state which resides in some subspace $\mathcal{H}^{(\alpha)} = \otimes_j \mathcal{H}_j^{(\alpha_j)}$, for some choice of $\alpha_1, \dots, \alpha_n$. Moreover, within the subspace $\mathcal{H}^{(\alpha)}$ the state can be written as a tensor product of two-particle states living in the spaces of the form $\mathcal{H}_{R,j}^{(\alpha_j)} \otimes \mathcal{H}_{L,j+1}^{(\alpha_{j+1})}$, tensored with some arbitrary single-particle states living in the $\mathcal{H}_{C,j}^{(\alpha_j)}$ spaces.

If the algorithm knows the correct choice of indices $\alpha_1, \dots, \alpha_n$, it can find such a ground state efficiently, as follows. Note that the descriptions of both the spaces $\mathcal{H}_j^{(\alpha_j)}$ and

²This can also be seen by analyzing the tree width of that tensor network and using the analysis of Ref. [17].

their divisions $\mathcal{H}_j^{(\alpha_j)} = \mathcal{H}_{L,j}^{(\alpha_j)} \otimes \mathcal{H}_{C,j}^{(\alpha_j)} \otimes \mathcal{H}_{R,j}^{(\alpha_j)}$ are derived from local properties of \mathcal{H}_j imposed by the two local Hamiltonians $H_{j-1,j}$ and $H_{j,j+1}$. The subdivision of \mathcal{H}_j in this way can be expressed as a solution to a set of quadratic homogeneous constraints. Since the dimension of \mathcal{H}_j and hence the number of variables is constant, it can be efficiently computed. If the algorithm knows the α_j 's, it therefore knows the description of the subspaces $\mathcal{H}_{L,j}^{(\alpha_j)} \otimes \mathcal{H}_{C,j}^{(\alpha_j)} \otimes \mathcal{H}_{R,j}^{(\alpha_j)}$ and the restriction of the $H_{j,j+1}$ to those spaces; it therefore just needs to find a ground state of linearly many two-particle Hamiltonians, which is an easy task. It is therefore enough for the algorithm to find the correct $\alpha_1, \dots, \alpha_n$ indices.

We do this using dynamic programming. The critical point in using dynamic programming here is that the energy contribution of $H_{\ell,\ell+1}$ depends only on the choice of α_ℓ and $\alpha_{\ell+1}$, so the choice of α_k for $k \leq j-1$ does not affect the energy of the $H_{\ell,\ell+1}$ terms for any $\ell \geq j$. Using this observation, the algorithm proceeds from left to right as follows. For the first term $H_{1,2}$, the algorithm finds the division into a direct sum of subspaces for particles 1 and 2. The algorithm keeps an optimal state (choice of α_1) and energy for each possible α_2 .

Then, in a general step, we assume at particle i that we have the following information for each index α_i : a list of indices $\alpha_1, \dots, \alpha_{i-1}$ such that the ground energy of the Hamiltonian of particles $1, \dots, i$ restricted to the subspaces $\mathcal{H}_1^{(\alpha_1)} \otimes \dots \otimes \mathcal{H}_i^{(\alpha_i)}$ is minimal. To continue to the next particle, we first compute the division into subspaces for particle $i+1$, indexed as α_{i+1} , and optimize for each subspace in turn. For each subspace, we consider all items in the previous list; for each item, we have a list of subspaces, one for each particle. We compute the minimal energy for each such restriction, including now the $H_{i,i+1}$ term in the calculation of the energy, restricted according to subspaces α_{i+1} and α_i , the last choice coming from the list. We pick the tail of the subspace of the $i+1$ particle to be the one which minimizes the terms up to that point.

Notice that in each step the dynamic program compares partial energies emerging from restricting the state to a different sector in the Hilbert space. These energies can be computed efficiently with polynomially many bits, namely up to exponentially good precision. Thus, this second algorithm achieves exact results for a somewhat larger set of Hamiltonians than our first algorithm, namely, those for which the partial energies will not be confused if the computations are done with exponentially good precision.

Note that even with this extremely good resolution, it might be the case that the ground energy is confused with a slightly excited energy which is, say, doubly exponentially close. We do not know of a good condition which would rule out the possibility of such very close energies, except for some very trivial assumptions such as requiring that all eigenvalues are integer numbers. For example, even if we require that the different entries in the terms in the Hamiltonian are all rationals smaller than 1 with denominator upper bounded by a constant, it is still not known how to rule out the possibility that two eigenvalues of the overall Hamiltonian are doubly exponentially close. This issue touches upon an open question in number theory related to sums of algebraic numbers; see the

open problem described in Ref. [21], which can be traced back to Ref. [22] (if not earlier), and also Ref. [23] and references therein.

ACKNOWLEDGMENTS

Dorit Aharonov is partially supported by ISF Grant Nos. 039-7549 and 039-8066, ARO Grant No. 030-7799, and SCALA Grant No. 030-7811. Sandy Irani is partially supported by NSF Grant No. CCF-0916181. Itai Arad acknowledges support by the ERC Starting Grant of Julia Kempe (PI). The main progress on the results reported on in this paper was made while the three authors were visiting the Erwin Schrödinger International Institute for Mathematical Physics (ESI) in Vienna, Austria.

APPENDIX: PROOF OF THE LEMMAS

Proof of Lemma 1. Let $\delta = v/72b$. We will occasionally use the assumption that $\delta \leq 1/72b$.

First we create a set $R(\delta)$ of real numbers in the interval $[0,1]$ such that for any real number in the range $[0,1]$, it is within δ of some element in $R(\delta)$. $R(\delta)$ will have $\lceil 1/2\delta \rceil$ elements. To create $R(\delta)$, we add $(2j+1)\delta$ for each integer j in the range from 0 through $\lceil 1/2\delta \rceil - 2$. Note that the largest point in $R(\delta)$ so far is in the range $[1-3\delta, 1-\delta)$. Then we add $1-\delta$ to $R(\delta)$.

Then using $R(\delta)$, we create a set $C(\delta)$ which is a set of complex scalars which form a net over all complex scalars with norm at most 1. Include $xe^{i2\pi y}$ for every $x, y \in R(\delta)$. There are $\lceil 1/2\delta \rceil^2 \leq (1/\delta)^2$ points in $C(\delta)$. For any complex number c if norm at most 1, there is a number c' in $C(\delta)$ such that $|c - c'| \leq 2\delta$.

To generate $S_{a,b}$, consider first the set S_1 of all possible $a \times b$ matrices with entries from $C(\delta)$. This set contains $|C(\delta)|^{ab}$ matrices. In the case where $a = 1$ and we only want entries with real, non-negative coefficients, we use $R(\delta)$ for the entries instead of $C(\delta)$ and the set contains $|R(\delta)|^b$ matrices (in fact, vectors). Then

1. remove any matrix from S_1 which has a row whose norm is greater than $1 + \sqrt{b}2\delta$ or less than $1 - \sqrt{b}2\delta$, to get S_2 ;
2. renormalize each row in every matrix in S_2 to get S_3 ;
3. remove any matrix from S_3 which has any two rows whose inner product is more than $9\sqrt{b}\delta$;
4. for any matrix in S_3 , apply the Gram-Schmidt procedure to the rows to form an orthonormal set.

We claim that the final set is the desired $S_{a,b}$. Note that the number of matrices is $O((1/\delta)^{2ab}) = O((72b/v)^{2ab})$, and the running time to produce the set is $O(a^2b(1/\delta)^{2ab}) = O(a^2b(72b/v)^{2ab})$, as required. What remains to show is that if A is any $a \times b$ matrix whose rows form an orthonormal set, then we can find a matrix B in $S_{a,b}$ which is close to it.

Let W be an $a \times b$ matrix. We will denote its i th row as W_i . Define the distance between two matrices $d(W, W')$ to be $\max_i \|W_i - W'_i\|$. Let X be the matrix obtained by rounding every entry in A to the nearest complex number in $C(\delta)$. Let Y be the matrix obtained after the rows of X are normalized and let Z be the matrix obtained after the rows of Y are transformed into an orthonormal set via the Gram-Schmidt procedure. We need to prove that $d(A, Z) \leq v$ and to show that $Z \in S_{a,b}$,

which would imply together that we can choose B in the lemma to be equal to Z .

We will now prove both of the above claims. For the second part we need to prove that X survives step 1 and Y survives step 3.

X survives step 1. Since each entry in $A - X$ has magnitude of at most 2δ , we know that $d(A, X) \leq \sqrt{b}2\delta$. In order to bound the norm of X_i , observe that

$$\sqrt{b}2\delta \geq \|A_i - X_i\| \geq \| \|A_i\| - \|X_i\| \|.$$

Since $\|A_i\| = 1$, it follows that $\|X_i\|$ lies in the range from $1 - \sqrt{b}2\delta$ to $1 + \sqrt{b}2\delta$ and it will survive step 1. We have

$$\begin{aligned} d(X, Y) &\leq \max_i \left\| A^i - \frac{1}{1 - \sqrt{b}2\delta} A^i \right\| \\ &= \frac{\sqrt{b}2\delta}{1 - \sqrt{b}2\delta} \leq \sqrt{b}2\delta(36/35). \end{aligned}$$

The latter inequality uses the assumption that $\delta \leq 1/72\sqrt{b}$. Using the triangle inequality for our distance $d(\cdot)$, we have that for any i $\|A_i - Y_i\| \leq (4 + \frac{2}{35})\sqrt{b}\delta$.

Y survives step 3. Now we need to bound the inner product of any two rows of Y in order to establish that it is not removed in step 3:

$$\begin{aligned} |\langle Y_i | Y_j \rangle| &= |\langle A_i + (Y_i - A_i) | A_j + (Y_j - A_j) \rangle| \\ &\leq |\langle A_i | A_j \rangle| + |\langle Y_i - A_i | Y_j - A_j \rangle| \\ &\quad + |\langle Y_i - A_i | A_j \rangle| + |\langle A_i | Y_j - A_j \rangle| \\ &\leq \|Y_i - A_i\| \|Y_j - A_j\| + \|Y_i - A_i\| \|A_j\| \\ &\quad + \|A_i\| \|Y_j - A_j\| \\ &\leq \left[\left(4 + \frac{2}{35}\right)^2 \sqrt{b}\delta + 2 \left(4 + \frac{2}{35}\right) \right] \sqrt{b}\delta \\ &\leq 9\sqrt{b}\delta. \end{aligned}$$

The second inequality uses the Chauchy-Schwartz inequality. The last inequality uses the fact that $\sqrt{b}\delta \leq 1/72$.

Bounding the distance $d(A, Z)$. Finally, we need to consider how much the matrix shifts as a result of the Gram-Schmidt procedure, to bound $d(Y, Z)$. Let $\mu = 9\sqrt{b}\delta = 9v/72\sqrt{b}$. Since $a \leq b$, by assumption in the lemma, we know that $\mu \leq 9v/72\sqrt{a}$. We use this latter bound in the next part of the proof since we are bounding quantities by a function of a instead of b . Since we assume that $v \leq 1/\sqrt{a}$, we can assume that $a\mu \leq 9/72$. Recall that the Gram-Schmidt procedure starts with $Z_1 = Y_1$. Then each Z_i is determined by first creating an un-normalized state \tilde{Z}_i :

$$\tilde{Z}_i = Y_i - \sum_{j=1}^{i-1} \langle Z_j | Y_i \rangle Z_j.$$

Then \tilde{Z}_i is normalized to 1. We will prove the following two properties by induction in i :

1. $|\langle Z_i | Y_j \rangle| \leq 2\mu$ for all j such that $j > i$,
2. $1 - 2\sqrt{a}\mu \leq \|\tilde{Z}_i\| \leq 1 + 2\sqrt{a}\mu$.

\tilde{Z}_1 is not defined, but we can take it to be Z_1 . The two properties clearly hold for Z_1 . Now by induction,

$$\begin{aligned} \|\tilde{Z}_i\| &= \left\| Y_i - \sum_{j=1}^{i-1} \langle Z_j | Y_i \rangle Z_j \right\| \\ &\leq \|Y_i\| + \left\| \sum_{j=1}^{i-1} \langle Z_j | Y_i \rangle Z_j \right\| \\ &= 1 + \left(\sum_{j=1}^{i-1} \langle Z_j | Y_i \rangle \langle Y_i | Z_j \rangle \right)^{1/2} \\ &\leq 1 + 2\sqrt{a}\mu. \end{aligned}$$

A similar argument can be used to show that $1 - 2\sqrt{a}\mu \leq \|\tilde{Z}_i\|$. Next we establish Property 1:

$$\begin{aligned} |\langle Y_k | Z_i \rangle| &= \frac{1}{\|\tilde{Z}_i\|} \left| \langle Y_k | Y_i \rangle - \sum_{j=1}^i \langle Z_j | Y_i \rangle \langle Y_k | Z_j \rangle \right| \\ &\leq \frac{1}{1 - 2\sqrt{a}\mu} \left[\mu + \sum_{j=1}^{i-1} 4\mu^2 \right] \\ &\leq \frac{\mu(1 + 4a\mu)}{1 - 2\sqrt{a}\mu} \leq 2\mu. \end{aligned}$$

The first inequality follows from the inductive hypothesis. The last inequality make use of the fact that $a\mu \leq 9/72$. Finally, to bound $\|Y_i - Z_i\|$, we have

$$\begin{aligned} \|Y_i - Z_i\| &\leq \left(1 - \frac{1}{\|\tilde{Z}_i\|}\right) \|Y_i\| + \frac{1}{\|\tilde{Z}_i\|} \left\| \sum_{j=1}^{i-1} \langle Z_j | Y_i \rangle Z_j \right\| \\ &\leq \frac{2\sqrt{a}\mu \|Y_i\|}{1 - 2\sqrt{a}\mu} + \frac{1}{1 - 2\sqrt{a}\mu} \left(\sum_{j=1}^{i-1} |\langle Z_j | Y_i \rangle|^2 \right)^{1/2} \\ &\leq \frac{4\sqrt{a}\mu}{1 - 2\sqrt{a}\mu} \leq 6\sqrt{a}\mu. \end{aligned}$$

The last inequality uses again the fact that $\sqrt{a}\mu \leq 9/72$. The total distance between A and Z is at most $5\sqrt{b}\delta + 6\sqrt{a}\mu$. Plugging in $\mu = 9\sqrt{b}\delta$ and using the fact that $a \leq b$, we get an upper bound of $59b\delta \leq v$ on the distance of A to Z , using the definition of δ . ■

Lemma 4. Let $|A\rangle, |B\rangle$ be two two-particle states and expand them in the standard basis of the first particle,

$$\begin{aligned} |A\rangle &= \sum_i a_i |i\rangle |A_i\rangle, \\ |B\rangle &= \sum_i b_i |i\rangle |B_i\rangle, \end{aligned}$$

such that $|A_i\rangle$ are normalized but not necessarily orthogonal to themselves and similarly the $|B_i\rangle$. Then

$$\|a - b\| = \left(\sum_i |a_i - b_i|^2 \right)^{1/2} \leq \|A - B\|. \quad (A1)$$

Proof.

$$\begin{aligned} \|a - b\|^2 &= \sum_i |a_i - b_i|^2 \\ &\leq \sum_i \|a_i|A_i\rangle - b_i|B_i\rangle\|^2 \\ &= \left\| \sum_i |i\rangle(a_i|A_i\rangle - b_i|B_i\rangle) \right\|^2 \\ &= \||A\rangle - |B\rangle\|^2. \end{aligned}$$

■

-
- [1] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
 [2] S. R. White, *Phys. Rev. B* **48**, 10345 (1993).
 [3] S. Östlund and S. Rommer, *Phys. Rev. Lett.* **75**, 3537 (1995).
 [4] S. Rommer and S. Östlund, *Phys. Rev. B* **55**, 2164 (1997).
 [5] I. Peschel, X. Wang, M. Kaulke, and K. H., eds., *Density-Matrix Renormalization—A New Numerical Method in Physics* (Springer-Verlag, Berlin, 1998), Lecture Notes in Physics, Vol. 528.
 [6] U. Schollwöck, *Rev. Mod. Phys.* **77**, 259 (2005).
 [7] A. Kitaev, A. Shen, and M. Vyalyi, *Classical and Quantum Computation* (AMS, Providence, RI, 2002).
 [8] S. Bravyi and M. Vyalyi, *Quantum Inf. Comput.* **5**(3), 187 (2005).
 [9] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, *SIAM J. Comput.* **37**, 166 (2007).
 [10] D. Aharonov, D. Gottesman, S. Irani, and J. Kempe, *Commun. Math. Phys.* **287**(1), 41 (2009).
 [11] D. Aharonov, I. Arad, Z. Landau, and U. Vazirani, in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing* (ACM, New York, NY, USA, 2009), p. 417.
 [12] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, *Quantum Inf. Comp.* **10**, 1 (2010).
 [13] N. Schuch, I. Cirac, and F. Verstraete, *Phys. Rev. Lett.* **100**, 250501 (2008).
 [14] J. Eisert, *Phys. Rev. Lett.* **97**, 260501 (2006).
 [15] M. B. Hastings, *J. Stat. Mech.* (2007) P08024.
 [16] N. Schuch and J. I. Cirac, *Phys. Rev. A* **82**, 012314 (2010).
 [17] I. Markov and Y. Shi, *SIAM J. Comput.* **38**, 963 (2008).
 [18] D. Aharonov, Z. Landau, and J. Makowsky, e-print [arXiv:quant-ph/0611156](http://arxiv.org/abs/quant-ph/0611156).
 [19] I. Arad and Z. Landau, *SIAM J. Comput.* **39**, 3089 (2010).
 [20] G. Vidal, *Phys. Rev. Lett.* **91**, 147902 (2003).
 [21] J. D. Demaine, J. S. B. Mitchell, and J. O’Rourke, *The Open Problems Project, Problem 33* [<http://maven.smith.edu/~orourke/TOPP/Welcome.html>].
 [22] J. O’Rourke, *Am. Math. Monthly* **88**, 769 (1981).
 [23] J. Qian and C. Wang, *Inf. Process. Lett.* **100**, 194 (2006).