Improved Approximation Algorithms for Bounded-Degree Local Hamiltonians

Anurag Anshu, David Gosset, Karen J. Morenz Korol, and Mehdi Soleimanifar

Phys. Rev. Lett. **127**, 250502 — Published 17 December 2021

DOI: 10.1103/PhysRevLett.127.250502

Improved approximation algorithms for bounded-degree local Hamiltonians

Anurag Anshu[1], David Gosset[2], Karen J. Morenz Korol[3], and Mehdi Soleimanifar[4]
[1] *Department of EECS & Challenge Institute for Quantum Computation, University of California, Berkeley, USA and Simons Institute for the Theory of Computing, Berkeley, California, USA.*
[2] *Department of Combinatorics and Optimization and Institute for Quantum Computing, University of Waterloo, Canada*
[3] *Department of Chemistry, University of Toronto, Canada and*
[4] *Center for Theoretical Physics, Massachusetts Institute of Technology, USA*

The low-temperature properties of interacting quantum systems are believed to require exponential resources to compute in the general case. Quantifying the extent to which such properties can be approximated using efficient algorithms remains a significant open challenge. Here we consider the task of approximating the ground state energy of two-local quantum Hamiltonians with bounded-degree interaction graphs. Most existing algorithms optimize the energy over the set of product states. We propose and analyze a family of shallow quantum circuits that can be used to improve the approximation ratio achieved by a given product state. The algorithm takes as input an n-qubit product state with variance Var and improves its energy by an amount proportional to Var^2/n. In a typical case, this results in an extensive improvement in the estimated energy. We extend our results to k-local Hamiltonians and entangled initial states.

Quantum computers are capable of efficiently computing the dynamics of quantum many-body systems [1], and it is anticipated that they can be useful for scientific applications in physics, materials science and quantum chemistry. The extent of the quantum advantage for other important simulation tasks, such as computing low temperature properties of quantum systems, is still unknown. In this paper we consider the task of approximating the ground state energy of local Hamiltonians. Here it is natural to expect some improvement over classical machines which cannot even store the state of such systems efficiently. Indeed, classical methods such as the mean-field or Hartree-Fock approximations do not capture the entanglement structure present in the true ground state.

Motivated by small quantum computers that may be available in the near future, there has been increased interest in devising algorithms that consume few quantum resources and can be implemented across a wide range of hardware platforms. In this vein, heuristic algorithms for ground state preparation have been proposed based on variationally minimizing the energy over the output states of shallow (low-depth) quantum circuits [2–4]. Although variational algorithms have been rigorously analyzed for specific problems and some limitations are known [5–8], no general treatment of their efficacy exists. Characterizing the advantage offered by shallow quantum circuits and variational quantum algorithms stands as a pressing challenge.

In this paper, we derive rigorous bounds on the performance of shallow quantum circuits in estimating the ground state energy of local Hamiltonians. For simplicity, we state our results for a system of qubits with two-local interactions. In the Supplemental Material, we discuss extensions of our results to k-local Hamiltonians.

To begin, let $G = (V, E)$ be a graph, and consider a Hamiltonian

$$H = \sum_{\{i,j\}\in E} h_{ij} \tag{1}$$

with $n = |V|$ qubits and nearest-neighbor interactions h_{ij} that act nontrivially only on qubits $\{i, j\}$ at vertices connected by an edge. We assume without loss of generality that $\|h_{ij}\| \leq 1$. We are interested in the problem of approximating the ground energy or smallest eigenvalue $\lambda_{\min}(H)$ of the Hamiltonian. It will be convenient to instead approximate the largest eigenvalue $\lambda_{\max}(H)$; this convention matches the one used in classical optimization and is without loss of generality, since $\lambda_{\min}(H) = -\lambda_{\max}(-H)$. In the worst case, the problem of estimating the largest eigenvalue $\lambda_{\max}(H)$ of Eq. (1) to within an additive error scaling inverse polynomially with n is believed to be intractable for quantum or classical computers [9]. Here we consider the approximation task where the goal is to compute an estimate $e \leq \lambda_{\max}(H)$ such that the approximation ratio $r \equiv e/\lambda_{max}(H)$ is as large as possible. We will also be interested in efficient quantum algorithms that prepare states $|\psi\rangle$ with good approximation ratios.

Besides describing local interactions encountered in physics, Hamiltonians of the form Eq. (1) can encode notable cost-functions considered in computer science and thus provide a physically motivated extension of the classical *approximation algorithm* setting [10]. For example, one may consider an Ising Hamiltonian for which $h_{ij} = (I - Z_iZ_j)/2$, where Z is the Pauli operator. This Hamiltonian is classical—that is, diagonal in the computational basis—and computing its maximum eigenvalue is equivalent to finding the Max-Cut of the graph G, a well-studied classical optimization problem. More generally, two-local quantum Hamiltonians may involve noncommuting terms such as Heisenberg interactions $h_{ij} = 1/4(I - X_iX_j - Y_iY_j - Z_iZ_j)$ (with Pauli X, Y and Z oper-

ators); the resulting optimization problem can be viewed as a quantum analogue of Max-Cut [11]. Quantum approximation algorithms aim to estimate the largest eigenvalue of such Hamiltonians and have been studied in several previous works. This includes the Heisenberg interactions mentioned above [11, 12] and more general settings in which the interaction terms h_{ij} are restricted to be positive semidefinite [13–15], or traceless [16, 17].

Despite considerable interest, the ultimate limits of efficient algorithms for quantum approximation algorithms are poorly understood. Approximation ratios approaching 1 are only known to be achievable for certain special families of graphs, including lattices or bounded-degree planar graphs using tensor product of $O(1)$-qubit states [18] or high degree graphs using tensor products of single-qubit states [13, 18, 19]. In certain cases, one may ascertain limitations on efficient achievable approximation ratios from the classical Probabilistically Checkable Proof (PCP) theorem [20–22], though stronger and more general limitations may follow from the quantum PCP conjecture if some version of it can be proven [23].

A quantum approximation algorithm typically outputs an estimate of the form $\langle v|H|v\rangle$ where $|v\rangle$ is a quantum state computed by the algorithm. A central challenge is to understand the structure of quantum states $|v\rangle$ that achieve high approximation ratios in the general case. Most existing algorithms are based on tensor products of one- or few-qubit states, while Ref. [12] also considers states prepared by shallow quantum circuits. In this work we describe conditions under which the performance of such algorithms can be *improved*. We restrict our attention to local Hamiltonians on bounded-degree graphs and consider an improvement strategy based on shallow quantum circuits.

Improvement of product states To this end, suppose we are given an n-qubit state $|v\rangle$ and a Hamiltonian Eq. (1) defined on a graph $G = (V, E)$ with maximum degree $d \geq 2$. It will be convenient to assume (without loss of generality) that G is d-regular—we can ensure this by possibly adding some local terms h_{ij} which are equal to zero. We imagine that $|v\rangle$ may be the output of some approximation algorithm such as the ones described above. Our aim is to efficiently compute a state with energy larger than $\langle v|H|v\rangle$. Moreover, we would like to increase this energy by an amount proportional to $|E|$ in order to guarantee that the *approximation ratio* is larger by some additive constant. We show that this is possible if the following two conditions hold:

(i) The variance of the energy, defined by

$$\mathrm{Var}_v(H) = \langle v|H^2|v\rangle - \langle v|H|v\rangle^2,$$

satisfies $\mathrm{Var}_v(H) = \Omega(|E|)$ [24].

(ii) The state $|v\rangle$ is a product state. That is, $|v\rangle = |v_1\rangle \otimes |v_2\rangle \otimes \ldots \otimes |v_n\rangle$ where each $|v_i\rangle$ is a single-qubit state.

More generally, the following theorem quantifies the energy increase as a function of $\mathrm{Var}_v(H)$ even if condition (i) is not satisfied.

Theorem 1. *Given a product state $|v\rangle$, we can efficiently compute a depth-$(d+1)$ quantum circuit U such that the state $|\psi\rangle = U|v\rangle$ satisfies*

$$\langle\psi|H|\psi\rangle \geq \langle v|H|v\rangle + \Omega\left(\frac{\mathrm{Var}_v(H)^2}{d^2|E|}\right). \tag{2}$$

This result applies broadly to quantum optimization problems, but does not provide any improvement when specialized to the classical setting. To see this, note that condition (i) is not satisfied in the purely classical case where $|v\rangle$ is a computational basis state and H is diagonal in the computational basis. Indeed, we have $\mathrm{Var}_v(H) = 0$ whenever $|v\rangle$ is an eigenstate of H. On the other hand, condition (i) is fairly mild in the quantum setting; in the Supplemental Material we show that it holds for a generic product state $|v\rangle$ whenever the Hamiltonian contains nontrivial interactions on each edge of the graph.

Simple examples demonstrate that neither of the two conditions alone is enough to even guarantee the existence of a state with approximation ratio better than $|v\rangle$ for large regular graphs. Condition (ii) alone is not sufficient because it is possible for a product state to have maximal energy $\lambda_{\max}(H)$ (i.e., this occurs for all classical Hamiltonians). To see that condition (i) is not sufficient, one can consider the Max-Cut Hamiltonian on (say) an even cycle graph, and let $|v\rangle$ be an equal superposition of two eigenstates of H, one with maximal energy $|E|$ and one with energy $|E| - \Theta(\sqrt{|E|})$. The resulting state has approximation ratio $1 - O(|E|^{-1/2})$ and variance $\mathrm{Var}_v(H) = \Omega(|E|)$. Thus condition (i) is satisfied, but the approximation ratio cannot be improved by an additive constant.

In the special case where $|v\rangle$ achieves the largest energy of any product state, we are able to strengthen the bound Eq. (2). We say that the product state $|v\rangle$ is *locally optimal* for H if for any single-qubit Pauli Q, we have

$$\frac{d}{d\phi}\langle v|e^{-i\phi Q}He^{i\phi Q}|v\rangle\big|_{\phi=0} = 0,$$

or equivalently $\langle v|[Q, H]|v\rangle = 0$. As we show in the Supplemental Material, the bound in Eq. (2) can be improved to $\langle v|H|v\rangle + \Omega(\frac{\mathrm{Var}_v(H)^2}{d|E|})$ for locally optimal states.

Generally, however, the improvement stated in Eq. (2) is optimal in the sense that there exists a Hamiltonian H and a product state $|v\rangle$ with $\mathrm{Var}_v(H) = \Theta(|E|)$ for

which

$$\lambda_{\max}(H) - \langle v|H|v\rangle \leq O\left(\frac{\mathrm{Var}_v(H)^2}{d^2|E|}\right). \tag{3}$$

For example, Eq. (3) is satisfied by the Hamiltonian with $h_{ij} = Z_i + Z_j$ on any d-regular graph and the product state $|v\rangle = (\cos(\theta)|0\rangle + \sin(\theta)|1\rangle)^{\otimes n}$, for any $\theta \in (0, \pi/2)$. In this simple case, the left-hand side can be computed exactly and is equal to $\frac{\mathrm{Var}_v(H)^2}{d^2|E|} \cdot \frac{\sin^2(\theta)}{\sin^4(2\theta)}$.

To establish Theorem 1, we consider a variational family of states obtained from $|v\rangle = \otimes_{i\in V}|v_i\rangle$ by applying a quantum circuit composed of nearest neighbor commuting gates on the interaction graph G. In particular, let $P_1, P_2, \ldots, P_n$ be any collection of single-qubit operators such that $\|P_i\| \leq 1$ and

$$\langle v_i|P_i|v_i\rangle = 0 \quad \text{for all} \quad i \in V.$$

Following [12], we define the circuit

$$V(\vec{\theta}) = \prod_{\{i,j\}\in E} e^{i\theta_{ij}P_iP_j} = e^{i\sum_{\{i,j\}\in E}\theta_{ij}P_iP_j}. \tag{4}$$

Here, $\vec{\theta}$ is an array of real parameters $\{\theta_{ij}\}_{\{i,j\}\in E}$. Since by assumption, the interaction graph G is d-regular, the quantum circuit $V(\vec{\theta})$ can be implemented with circuit depth $d+1$. It is not hard to see that this variational family includes as a special case the level-1 Quantum Approximate Optimization Algorithm (QAOA) for 2-local classical Hamiltonians [3]. For a given choice of operators $\{P_i\}_{i\in V}$, the following theorem lower bounds the improvement in the energy after applying the the quantum circuit $V(\vec{\theta})$ to $|v\rangle$.

Theorem 2. *Let $|v\rangle$ be a product state and $|\psi\rangle = V(\vec{\theta})|v\rangle$ be the state prepared by the quantum circuit Eq. (4). Define the positive real parameter α by*

$$\alpha = \mathbb{E}_{\{i,j\}\in E}|\langle v_i, v_j|[P_iP_j, h_{ij}]|v_i, v_j\rangle|, \tag{5}$$

where the expectation is with respect to the uniform distribution over the edges. There is an efficient classical algorithm to select parameters $\vec{\theta}$ satisfying

$$\langle\psi|H|\psi\rangle \geq \langle v|H|v\rangle + \Omega\big(|E|\alpha^2/d\big). \tag{6}$$

Proof. Write N_{ij} for the set of edges $\{k,\ell\} \in E$ incident to a given edge $\{i,j\} \in E$. The latter edge is included as well, i.e., $\{i,j\} \in N_{ij}$. Consider the energy of a term

$$\langle\psi|h_{ij}|\psi\rangle = \langle v|V(\vec{\theta})^\dagger h_{ij} V(\vec{\theta})|v\rangle.$$

The gates in $V(\vec{\theta})$ which are associated with edges that are not incident with $\{i,j\}$ can be cancelled, leaving $\langle v|V_{ij}^\dagger h_{ij} V_{ij}|v\rangle$ where $V_{ij} = \prod_{\{k,\ell\}\in N_{ij}} e^{i\theta_{k\ell}P_kP_\ell}$. Thus

$$\langle\psi|h_{ij}|\psi\rangle = \langle v|h_{ij}|v\rangle + \sum_{m=1}^{\infty}\frac{i^m}{m!}\langle v|\Bigg[\sum_{\{k,\ell\}\in N_{ij}} -\theta_{k\ell}P_kP_\ell, h_{ij}\Bigg]_m|v\rangle. \tag{7}$$

Here, $[A,B]_m$ is the m-nested commutator $[A,[A,\ldots[A,B]]]$. Using the fact that $\langle v_k|P_k|v_k\rangle = 0$ for all k, the $m=1$ term simplifies to

$$\sum_{\{k,\ell\}\in N_{ij}} -i\theta_{k\ell}\langle v|\left[P_kP_\ell, h_{ij}\right]|v\rangle = -i\theta_{ij}\langle v|[P_iP_j, h_{ij}]|v\rangle. \tag{8}$$

At this stage, we make the choice

$$\theta_{ij} = \theta \cdot \mathrm{sign}\left(-i\langle v|[P_iP_j, h_{ij}]|v\rangle\right), \tag{9}$$

where the parameter θ will be determined later. Substituting in Eq. (8) gives

$$\sum_{\{k,\ell\}\in N_{ij}} -i\theta_{k\ell}\langle v|\left[P_kP_\ell, h_{ij}\right]|v\rangle = \theta|\langle v_i, v_j|[P_iP_j, h_{ij}]|v_i, v_j\rangle|. \tag{10}$$

For $m > 1$, we have

$$\left|\langle v|\Bigg[\sum_{\{k,\ell\}\in N_{ij}} -\theta_{k\ell}P_kP_\ell, h_{ij}\Bigg]_m|v\rangle\right| \leq \sum_{\substack{\{k_1,\ell_1\},\{k_2,\ell_2\},\ldots\\ \{k_m,\ell_m\}\in N_{ij}}} \theta^m\left|\langle v|\left[P_{k_1}P_{\ell_1}, [\ldots,[P_{k_m}P_{\ell_m}, h_{ij}]]\right]|v\rangle\right|.$$

The only nonzero terms are those in which the expression $\langle v_s|P_s|v_s\rangle$ does not appear. To upper bound the number of nonzero terms, we count the number of tuples $(\{k_1,\ell_1\},\{k_2,\ell_2\},\ldots\{k_m,\ell_m\})$ such that no vertex in $V\setminus\{i,j\}$ appears exactly once. An upper bound is as follows (a proof is provided in the Supplemental Material).

Claim 1. *Let $m \geq 2$. The number of ordered tuples of edges $(\{k_1,\ell_1\},\{k_2,\ell_2),\ldots,\{k_m,\ell_m\}) \in N_{ij}^{\times m}$ in which no vertex in $V\setminus\{i,j\}$ appears exactly once is at most $(2m\sqrt{d})^m$.*

Finally, using Eq. (9) and the fact that $\|h_{ij}\|, \|P_i\| \leq 1$, we can upper bound

$$\theta^m\left|\langle v|\left[P_{k_1}P_{\ell_1}, [\ldots,[P_{k_m}P_{\ell_m}, h_{ij}]]\right]|v\rangle\right| \leq (2\theta)^m.$$

Thus, the sum of all $m > 1$ terms in Eq. (7) has magnitude at most

$$\sum_{m=2}^{\infty}\frac{1}{m!}\left(4m\sqrt{d}\right)^m\theta^m \leq \sum_{m=0}^{\infty}\left(4e\sqrt{d}\theta\right)^{m+2} \leq 32e^2d\theta^2$$

assuming $\theta \leq \frac{1}{8e\sqrt{d}}$ (where we used the bound $m^m/m! \leq e^m$). Combining with Eqs. (7,10) and summing over all $\{i,j\} \in E$, we get

$$\langle\psi|H|\psi\rangle \geq \langle v|H|v\rangle + |E|\left(\theta\alpha - 32e^2d\theta^2\right).$$

Choosing $\theta = O(\alpha/d)$ gives the desired lower bound. ∎

Let us now see how Theorem 1 is obtained as a consequence of Theorem 2. The lower bound (6) applies to any choice of operators $\{P_i\}_{i\in V}$. We will choose these operators in a way that gives the variance bound Eq. (2). In the following, for convenience and without loss of generality, we shall work in a local basis in which our initial product state is $|v\rangle \equiv |0^n\rangle$. Our starting point is the observation that the variance of a 2-local Hamiltonian can be expressed in this basis as

$$\mathrm{Var}_v(H) = \langle 0^n|HQ_1H|0^n\rangle + \langle 0^n|HQ_2H|0^n\rangle,$$

where Q_t is the projector onto computational basis states with Hamming weight $t \in \{1,2\}$. This implies that

$$\langle 0^n|HQ_tH|0^n\rangle \geq \mathrm{Var}_v(H)/2 \tag{11}$$

for some $t \in \{1,2\}$. Suppose $t = 2$ and let X_i,Y_i, and Z_i be the Pauli operators. We define α_1 to be the RHS of Eq. (5) with $P_i = X_i$ for all i, and similarly α_2 with $P_i = (X_i + Y_i)/\sqrt{2}$ for all i. By a direct calculation,

$$\begin{aligned}\alpha_1 &= \frac{2}{|E|}\sum_{\{i,j\}\in E} |\mathrm{Im}\left(\langle 11|h_{ij}|00\rangle\right)| \\ \alpha_2 &= \frac{2}{|E|}\sum_{\{i,j\}\in E} |\mathrm{Re}\left(\langle 11|h_{ij}|00\rangle\right)| \end{aligned} \tag{12}$$

and therefore

$$\langle 0^n|HQ_2H|0^n\rangle = \sum_{\{i,j\}\in E} |\langle 11|h_{ij}|00\rangle|^2 \leq |E|\left(\frac{\alpha_1+\alpha_2}{2}\right).$$

This means $\max\{\alpha_1,\alpha_2\} \geq |E|^{-1}\langle 0^n|HQ_2H|0^n\rangle$ which together with Eq. (11) implies that when $t = 2$, we can efficiently find a series of operators P_i such that the parameter α satisfies $\alpha \geq (2|E|)^{-1}\mathrm{Var}_v(H)$. By plugging this in Eq. (6), we obtain $\langle\psi|H|\psi\rangle \geq \langle v|H|v\rangle + \Omega(\frac{\mathrm{Var}_v(H)^2}{d|E|})$. Thus if $t = 2$ we obtain a better lower bound than the one claimed in Theorem 1. Otherwise, if $t = 1$, then a simple calculation (reproduced in the Supplemental Material) shows that one can efficiently compute a product state with energy at least $\langle v|H|v\rangle + \Omega(\frac{\mathrm{Var}_v(H)^2}{d^2|E|})$. In general, the choice between $t = 1$ and $t = 2$ can be efficiently determined. Thus we obtain Theorem 1. In the Supplemental Material, we show that if $|v\rangle$ is locally optimal for H, then $\langle 0^n|HQ_1H|0^n\rangle = 0$ and $t = 2$, so we obtain the better bound described above.

Let us briefly illustrate how these results can be applied to the quantum Max-Cut Hamiltonian considered in Refs. [11, 12]. The Hamiltonian is built from local terms $h_{ij} = w_{ij}\Pi_{ij}$, where $0 \leq w_{ij} \leq 1$ and $\Pi_{ij} = (I - X_iX_j - Y_iY_j - Z_iZ_j)/4$ is the projector onto the antisymmetric state of two qubits. This Hamiltonian has the special feature that any product state $|v\rangle$ is locally optimal, and moreover, we have $|\langle v_i^\perp, v_j^\perp|h_{ij}|v_i, v_j\rangle| = \langle v_i, v_j|h_{ij}|v_i, v_j\rangle$ for $|v^\perp\rangle$ orthogonal to $|v\rangle$. Therefore

$$\mathrm{Var}_v(H) = \sum_{\{i,j\}\in E}\langle v|h_{ij}|v\rangle^2 \geq |E|^{-1}\langle v|H|v\rangle^2$$

using Cauchy-Schwarz. We may then efficiently compute a state $|\psi\rangle$ such that

$$\langle\psi|H|\psi\rangle \geq \langle v|H|v\rangle + \Omega\left(\frac{\langle v|H|v\rangle^4}{d|E|^3}\right). \tag{13}$$

We see that if the initial state has approximation ratio $\langle v|H|v\rangle/|E| = r$ then the state $|\psi\rangle$ improves this to $r + \Omega(r^4/d)$ [25].

This example demonstrates the power of Theorem 1 and shows that for the quantum Max-Cut problem, the approximation ratio of *any* product state can be improved by applying a shallow quantum circuit. For more general two-local Hamiltonians, we can guarantee an improvement in the approximation ratio whenever the condition $\mathrm{Var}_v(H) = \Omega(|E|)$ holds, which we expect for typical product states and Hamiltonians. Below we discuss two natural extensions of our results. First, we ask whether one can improve approximation ratios attained by more general families of quantum states. Along these lines, we provide an extension of Theorem 1 to the more general case where $|v\rangle$ is any state prepared by a quantum circuit of depth $D = O(1)$. Next, we show how one can improve the approximation ratio achieved by a random product state $|v\rangle$. Using Theorem 2, we show that the approximation ratio can be improved by $\Omega(1/d)$ for any Hamiltonian with nontrivial two-local interactions, and by $\Omega(1/\sqrt{d})$ if the interaction graph is triangle-free.

Improvement of bounded-depth states Recall that for any n-qubit quantum circuit and any qubit $j \in [n]$, we may define the lightcone $\mathcal{L}(j) \subseteq [n]$ which consists of all output qubits that are causally connected to j. Define the maximum lightcone size $\ell = \max_{j\in[n]}\mathcal{L}(j)$. We have $\ell \leq 2^D$ for any depth D circuit composed of two-qubit gates.

Theorem 3. *Let $|v\rangle = W|0^n\rangle$ where W is a quantum circuit with maximum lightcone size ℓ. There is an efficient classical algorithm that computes a quantum circuit U such that $|\psi\rangle = U|v\rangle$ satisfies*

$$\langle\psi|H|\psi\rangle = \langle v|H|v\rangle + \Omega\left(\frac{\mathrm{Var}_v(H)^2}{\ell^{10}d^2|E|}\right).$$

For constant-depth circuits we have $\ell = O(1)$ and we get the same asymptotic energy improvement as we established previously in Theorem 1 for product states. However, here the circuit U is not constant-depth. In the Supplemental Material, we show that the improvement stated above can also be obtained for states $|v\rangle$ that are the unique ground states of a gapped ℓ-local Hamiltonian F. Thus, Theorem 3 extends to a broad

class of tensor network states (such as PEPS of low bond dimension) with a gapped parent Hamiltonian.

The theorem provides limitations on the energy that can be achieved by any state $|v\rangle$ produced by a bounded-depth circuit. In particular, since $\langle\psi|H|\psi\rangle \leq \lambda_{\max}(H)$, we find that

$$\langle v|H|v\rangle \leq \lambda_{\max}(H) - \Omega\left(\frac{\mathrm{Var}_v(H)^2}{\ell^{10}d^2|E|}\right).$$

This shows that the approximation ratio achievable by constant-depth states $|v\rangle$ with $\mathrm{Var}_v(H) = \Omega(|E|)$ is bounded away from 1. An interesting direction for future work is to explore whether one can use this fact to exhibit new local Hamiltonian systems with the almost-linear NLTS (No Low-energy Trivial States) property [26, 27].

Improvement of random assignments Given an instance of a (classical) constraint satisfaction problem, one may consider the trivial algorithm in which each variable is chosen independently and uniformly at random. Remarkably, efficient algorithms which improve over the approximation ratio achieved by this simple strategy are not likely to exist in the general case [28]. On the other hand, for structured cases such as bounded-degree graphs, improvement is possible. In particular, on degree-d graphs, one can efficiently find an assignment satisfying a $\mu + \Omega(\frac{1}{d})$ fraction of constraints [29]. Here μ is the expected fraction of constraints satisfied by a uniformly random assignment. It has been shown that when a degree-d graph is triangle-free, there are efficient "local" algorithms that find a binary string satisfying a $\mu + \Omega(\frac{1}{\sqrt{d}})$ fraction of constraints by starting with a uniformly random assignment [30, 31] or quantum superposition [32] and then locally updating each bit/qubit as a function of the state of its neighbors.

Below we show that this optimal dependence on d can be recovered and generalized to the local Hamiltonian setting by applying our algorithm in Theorem 2 to a randomly chosen product state. For randomly chosen $|v\rangle$, the parameter α in Theorem 2 can be related to the 2-norm of the quadratic terms in the Pauli expansion of the Hamiltonian. More precisely, for an n-qubit operator $O = \sum_{i<j}\sum_{x,y} f_{xy}^{ij}\sigma_x^i \otimes \sigma_y^j$ where $\sigma_0 = I$ and $\{\sigma_1, \sigma_2, \sigma_3\}$ are the Pauli matrices, we define

$$\mathrm{quad}(O) = \sum_{i<j}\sum_{x>0,y>0}(f_{xy}^{ij})^2.$$

Theorem 4. *There is an efficient randomized algorithm which computes a depth-$(d+1)$ circuit U such that $|\psi\rangle = U|v\rangle$ satisfies*

$$\mathbb{E}_v\langle\psi|H|\psi\rangle \geq \mathbb{E}_v\langle v|H|v\rangle + \Omega\left(\frac{\mathrm{quad}(H)^2}{d|E|}\right).$$

If the graph is triangle-free then the right-hand side can be replaced with $\mathbb{E}_v\langle v|H|v\rangle + \Omega\left(\frac{\mathrm{quad}(H)}{\sqrt{d}}\right)$.

The proof of Theorem 4 is provided in the Supplemental Material. We also show that for triangle-free graphs one can efficiently compute *product* states matching the approximation ratios quoted above using a local classical algorithm similar to the ones from Refs. [30, 31]. Thus, low depth quantum circuits are not necessary to achieve the $\Omega(1/\sqrt{d})$ scaling; see also Ref. [16] which establishes similar bounds for the general (not necessarily triangle-free) case. Nevertheless, one may take the output product state of such algorithms and improve it further using the shallow quantum circuit from Theorem 1.

Discussion For local Hamiltonian problems on bounded-degree graphs, we showed that the approximation ratio achieved by a product state can be improved by a shallow quantum circuit, assuming a mild condition on its variance. Our quantum algorithm generalizes the level-1 Quantum Approximate Optimization Algorithm (QAOA) and extends its applicability beyond classical cost functions. By applying our algorithm to randomly chosen product states we generalized known algorithms for bounded-occurrence classical constraint satisfaction problems. Our results quantify the improvement that shallow quantum circuits can provide over methods based on product states.

Acknowledgments AA acknowledges support from the NSF QLCI program through grant number OMA-2016245. DG acknowledges the support of the Natural Sciences and Engineering Research Council of Canada through grant number RGPIN-2019-04198, the Canadian Institute for Advanced Research, and IBM Research. KJMK acknowledges support from NSERC Vanier Canada Graduate Scholarship. MS was supported by NSF grant CCF-1729369, a Samsung Advanced Institute of Technology Global Research Cluster and grant number FXQi-RFP-1811A from the Foundational Questions Institute and Fetzer Franklin Fund, a donor advised fund of Silicon Valley Community Foundation.

[1] Seth Lloyd, "Universal quantum simulators," Science , 1073–1078 (1996).
[2] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'brien, "A variational eigenvalue solver on a photonic quantum processor," Nature communications **5**, 4213 (2014).
[3] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann, "A quantum approximate optimization algorithm," arXiv preprint arXiv:1411.4028 (2014).

[4] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," Nature **549**, 242–246 (2017).

[5] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven, "Barren plateaus in quantum neural network training landscapes," Nature communications **9**, 1–6 (2018).

[6] Edward Farhi, David Gamarnik, and Sam Gutmann, "The quantum approximate optimization algorithm needs to see the whole graph: A typical case," (2020), arXiv:2004.09002 [quant-ph].

[7] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang, "Obstacles to variational quantum optimization from symmetry protection," Phys. Rev. Lett. **125**, 260505 (2020).

[8] Sergey Bravyi, David Gosset, and Ramis Movassagh, "Classical algorithms for quantum mean values," Nature Physics **17**, 337–341 (2021).

[9] In particular, a decision version of this problem is complete for the complexity class QMA which is a quantum generalization of NP.

[10] Vijay V Vazirani, *Approximation algorithms* (Springer Science & Business Media, 2013).

[11] Sevag Gharibian and Ojas Parekh, "Almost optimal classical approximation algorithms for a quantum generalization of Max-Cut," in *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*, LIPIcs. Leibniz Int. Proc. Inform., Vol. 145 (Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019) pp. Art. No. 31, 17.

[12] Anurag Anshu, David Gosset, and Karen Morenz, "Beyond product state approximations for a quantum analogue of max cut," in *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)* (Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020).

[13] Sevag Gharibian and Julia Kempe, "Approximation algorithms for qma-complete problems," SIAM Journal on Computing **41**, 1028–1050 (2012).

[14] Sean Hallgren, Eunou Lee, and Ojas Parekh, "An approximation algorithm for the max-2-local hamiltonian problem," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2020)* (Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020).

[15] Ojas Parekh and Kevin Thompson, "Beating random assignment for approximating quantum 2-local hamiltonian problems," (2020), arXiv:2012.12347 [quant-ph].

[16] Aram W Harrow and Ashley Montanaro, "Extremal eigenvalues of local hamiltonians," Quantum **1**, 6 (2017).

[17] Sergey Bravyi, David Gosset, Robert König, and Kristan Temme, "Approximation algorithms for quantum many-body problems," Journal of Mathematical Physics **60**, 032203 (2019).

[18] Nikhil Bansal, Sergey Bravyi, and Barbara M Terhal, "Classical approximation schemes for the ground-state energy of quantum and classical ising spin hamiltonians on planar graphs," arXiv preprint arXiv:0705.1115 (2007).

[19] Fernando GSL Brandao and Aram W Harrow, "Product-state approximations to quantum states," Communications in Mathematical Physics **342**, 47–80 (2016).

[20] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy, "Proof verification and the hardness of approximation problems," J. ACM **45**, 501–555 (1998).

[21] Sanjeev Arora and Shmuel Safra, "Probabilistic checking of proofs: a new characterization of NP," J. ACM **45**, 70–122 (1998).

[22] Irit Dinur, "The PCP theorem by gap amplification," J. ACM **54**, Art. 12, 44 (2007).

[23] Dorit Aharonov, Itai Arad, and Thomas Vidick, "Guest column: the quantum pcp conjecture," Acm sigact news **44**, 47–79 (2013).

[24] I.e. there is a universal constant c, such that asymptotically $\mathrm{Var}_v(H) \geq c \cdot |E|$.

[25] A better bound can be obtained by directly computing the parameter α for a randomized choice of operators $\{P_i\}$. In that case, $\mathbb{E}\alpha \geq \Omega(\langle v|H|v\rangle/|E|)$ which results in an improvement of $\Omega\left(\langle v|H|v\rangle^2/(d|E|)\right)$.

[26] Anurag Anshu and Chinmay Nirkhe, "Circuit lower bounds for low-energy states of quantum code hamiltonians," (2021), arXiv 2011.02044.

[27] Michael H. Freedman and Matthew B. Hastings, "Quantum systems on non-k-hyperfinite complexes: A generalization of classical statistical mechanics on expander graphs," Quantum Info. Comput. **14**, 144–180 (2014).

[28] Johan Haastad, "Some optimal inapproximability results," J. ACM **48**, 798–859 (2001).

[29] Johan Haastad, "On bounded occurrence constraint satisfaction," Inform. Process. Lett. **74**, 1–6 (2000).

[30] Boaz Barak, Ankur Moitra, Ryan O'Donnell, Prasad Raghavendra, Oded Regev, David Steurer, Luca Trevisan, Aravindan Vijayaraghavan, David Witmer, and John Wright, "Beating the random assignment on constraint satisfaction problems of bounded degree," (2015), arXiv:1505.03424 [cs.CC].

[31] Matthew B Hastings, "Classical and quantum bounded depth approximation algorithms," arXiv preprint arXiv:1905.07047 (2019).

[32] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann, "A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem," (2015), arXiv:1412.6062 [quant-ph].

[33] Ryan O'Donnell, *Analysis of boolean functions* (Cambridge University Press, 2014).