Beating the Natural Grover Bound for Low-Energy Estimation and State Preparation

Harry Buhrman,[1,2,3,*] Sevag Gharibian,[4,†] Zeph Landau,[5,‡] François Le Gall,[6,§] Norbert Schuch,[7,8,‖] and Suguru Tamaki[9,¶]

[1]*Quantinuum, Terrington House 13-15 Hills Road Cambridge CB2 1NL, United Kingdom*
[2]*University of Amsterdam, Amsterdam, The Netherlands*
[3]*QuSoft, Science Park 123, 1098 XG Amsterdam, The Netherlands*
[4]*Paderborn University, Department of Computer Science and Institute for Photonic Quantum Systems (PhoQS), Warburger Strasse 100, 33098 Paderborn, Germany*
[5]*University of California, Berkeley, Department of Computer Science, Berkeley, California 94706, USA*
[6]*Nagoya University, Graduate School of Mathematics, Furocho, Chikusaku, 464-8602 Nagoya, Japan*
[7]*University of Vienna, Faculty of Mathematics, Oskar-Morgenstern-Platz 1, 1090 Wien, Austria*
[8]*University of Vienna, Faculty of Physics, Boltzmanngasse 5, 1090 Wien, Austria*
[9]*University of Hyogo, Graduate School of Information Science, 8-2-1 Gakuennishi-machi, Nishi-ku, 651-2197 Kobe, Japan*

(Received 4 July 2024; revised 4 February 2025; accepted 28 May 2025; published 15 July 2025)

Estimating ground state energies of many-body Hamiltonians is a central task in many areas of quantum physics. In this Letter, we give quantum algorithms which, given any k-body Hamiltonian H, compute an estimate for the ground state energy and prepare a quantum state achieving said energy, respectively. Specifically, for any $\varepsilon > 0$, our algorithms return, with high probability, an estimate of the ground state energy of H within additive error εM, or a quantum state with the corresponding energy. Here, M is the total strength of all interaction terms, which in general is extensive in the system size. Our approach makes no assumptions about the geometry or spatial locality of interaction terms of the input Hamiltonian and thus handles even long-range or all-to-all interactions, such as in quantum chemistry, where lattice-based techniques break down. In this fully general setting, the run-time of our algorithms scales as $2^{cn/2}$ for $c < 1$, yielding the first quantum algorithms for low-energy estimation breaking a standard square root Grover speedup for unstructured search. The core of our approach is remarkably simple, and relies on showing that an extensive fraction of the interactions can be neglected with a controlled error. What this ultimately implies is that even arbitrary k-local Hamiltonians have structure in their low energy space, in the form of an exponential-dimensional low energy subspace.

DOI: 10.1103/29qw-bssx

Determining the properties of a quantum system at low energies is one of the central problems in the study of complex quantum many-body systems. It is most challenging in settings with long-range or all-to-all interactions, encountered for instance in the study of molecules in quantum chemistry [1,2], but also in a variety of cold atomic and molecular as well as solid state systems [3–9]. At the same time, all-to-all interactions also appear naturally in quantum complexity theory when encoding computationally hard tasks into quantum Hamiltonians [10,11], and would necessarily [12] have to be present in the sought-for quantum PCP theorem [13].

The ability of both classical and quantum algorithms to compute low-energy properties of quantum many-body systems is fundamentally limited by quantum complexity theory, which poses restrictions on what generally applicable algorithms can provably achieve. Finding the ground state energy of a system of n spins up to precision $1/\mathrm{poly}(n)$ is known to be QMA-hard (where QMA is the quantum version of NP) [10,11,14–16], and thus hard even for quantum computers.

To make the problem more tractable, one can aim for some given *extensive accuracy* in terms of the total energy scale. If the total strength of the interactions scales like the system size, as expected in physical systems, this amounts to a constant (but arbitrarily small) energy density, and thus allows the study of properties of the system which are

*Contact author: hbuhrman@gmail.com
†Contact author: sevag.gharibian@upb.de
‡Contact author: zeph.landau@gmail.com
§Contact author: legall@math.nagoya-u.ac.jp
‖Contact author: norbert.schuch@univie.ac.at
¶Contact author: suguru.tamaki@gmail.com

stable at finite temperature. For systems with spatially local interactions, e.g., on a lattice, this problem can be solved efficiently (for a fixed precision) by cutting the system into constant-sized patches and neglecting the couplings between those patches [12]. However, this is no longer possible for sytems with all-to-all interactions. In fact, the problem must remain at least NP-hard, as the PCP theorem implies that constraint satisfaction problems (which map to Hamiltonians with long-range couplings) remain NP-hard even for determining whether a fixed fraction of terms is violated [17,18]; the validity of a potential quantum PCP theorem [13] would even imply its QMA-hardness.

In the light of these hardness results and the unstructured nature of the problem, the best run-time we can expect for a classical algorithm for estimating properties at constant energy density is exponential, by exact diagonalization methods such as the Lanczos algorithm [19,20]. In fact, the strong exponential time hypothesis states that classical algorithms *must* run essentially in time scaling with 2^n [21]. On a quantum computer, we can use Grover search, leading to a quadratic speedup and thus a runtime of $O(2^{n/2})$ for qubit systems [22–25]. These methods can be improved if we are given access to a *trial* or *guiding* state for the low-energy subspace, as in recent quantum algorithms for ground state energy estimation [26,27]; however, the computation of such guiding states is a serious bottleneck, e.g., in current quantum chemistry applications [28].

In this Letter, we provide a quantum algorithm which provides a superquadratic speedup for computing properties of a general Hamiltonian with all-to-all interactions and k-body terms at fixed energy density. More specifically, we show that given a Hamiltonian $H = \sum h_\alpha$ on $(\mathbb{C}^d)^{\otimes n}$ with arbitrary k-body terms, but with no restriction on the geometry of the interactions, there exists a quantum algorithm with run-time $O(d^{(1-c_\varepsilon)n/2}/\varepsilon)$, with $c_\varepsilon = (\varepsilon/k+\varepsilon) > 0$, which returns an estimate of the ground state energy up to any extensive error $\varepsilon \sum \|h_\alpha\|$, as well as a quantum algorithm which outputs a state within that energy range, from which subsequently properties stable at low-energy density can be computed. At the same time, our results also yield classical algorithms for the same task with quadratically larger run-time, which therefore in turn outperform exact diagonalization.

The key property which we prove and on which our results are built is that even for systems with nonlocal interactions, an extensive fraction of the interactions can be neglected with a controlled error, just as for spatially local systems. This in particular also implies that the number of states below a certain energy grows exponentially already at low energies, again just as for spatially local Hamiltonians—differently speaking, even all-to-all interactions are limited in the extent to which they can rule out low-energy states. This allows us to construct a nontrivial guiding state or, alternatively, to restrict the search space, and this way gives rise to algorithms beating the natural bound based on Grover search.

Setting and summary of results—Let us now summarize our results in more detail. We consider n d-level systems $(\mathbb{C}^d)^{\otimes n}$ which interact via a Hamiltonian,

$$H = \sum_{\alpha=1}^{m} h_\alpha, \tag{1}$$

where each Hamiltonian term h_α is a k-body term, i.e., it acts on at most k of those systems, but it does not obey any spatial (geometric) locality (hence, $m \leq n^k$). We will adopt the terminology "k-local" for such interactions in the following, as is standard in quantum complexity theory. For simplicity, we will restrict to the case of qubits, $d = 2$. Define $M \coloneqq \sum_{\alpha=1}^{m} \|h_\alpha\|$. Let $E_0(H)$ denote the ground state energy of H. We will be interested in the task of determining $E_0(H)$ up to multiplicative accuracy ε, i.e., to find an $\hat{E}$ with

$$|\hat{E} - E_0(H)| \leq \varepsilon M, \tag{2}$$

or to prepare a state with energy in that range. Note that for physical systems with geometrically local or sufficiently rapidly decaying and uniform interactions, M will be bounded by a constant times the system size n, and thus (2) amounts to an extensive accuracy, while for a setting with $\|h_\alpha\| \leq 1$, as natural for complexity theoretic scenarios, $M \leq m$. In the following, the only assumption on M which we will make is that it is upper bounded by a polynomial in n, which in particular is satisfied in the aforementioned scenarios.

Our first result is the construction of a new Hamiltonian H' which acts on a smaller number n' of spins, and whose ground state energy $E_0(H')$ satisfies that $|E_0(H) - E_0(H')| \leq \delta M$, where δ can be tuned by choosing n'; specifically, H' is obtained from H by dropping the spins with the weakest couplings. This implies that any algorithm which determines $E_0(H')$ to sufficient accuracy also gives an approximation to $E_0(H)$ up to extensive accuracy. In particular, by suitably choosing n', this allows us to show the following (here O^* denotes scaling up to a prefactor polynomial in n):

Theorem 1—Let $\varepsilon > 0$. There exists a quantum algorithm computing an $\hat{E}$ satisfying $|\hat{E} - E_0(H)| \leq \varepsilon M$ with high probability in time

$$O^*\left(2^{(1-\frac{\varepsilon}{k+\varepsilon})n/2}\varepsilon^{-1}\right). \tag{3}$$

There also exists a classical algorithm for this task with run-time $O^*(2^{[1-(\varepsilon/k+\varepsilon)]n}\varepsilon^{-1})$.

In fact, we are able to show a stronger connection: given a low-energy state of H', we can construct/prepare (in a systematic and efficient way) low-energy states of H. This has two consequences. First, it gives us an algorithm to prepare a low-energy state of H in sub-Grover time, by first preparing a low-energy state of H' and then transforming it to a low-energy state of H. Specifically, we show the following:

Theorem 2—Let $\varepsilon > 0$. There exists a quantum algorithm preparing, with high probability, a mixed state ρ with energy $\mathrm{Tr}[H\rho] \le E_0(H) + \varepsilon M$ in time

$$O^*\left(2^{(1-\frac{\varepsilon}{2k+\varepsilon})n/2}\varepsilon^{-1}\right). \tag{4}$$

As a second consequence, the connection between the low-energy space of H' and H allows us to construct an entire low-energy subspace whose dimension grows exponentially with energy. Specifically, we find that for any ε, the number $\mathcal{C}(H,\varepsilon)$ of eigenstates of H with energy below $E_0(H) + \varepsilon M$ scales as

$$\mathcal{C}(H,\varepsilon) \ge 2^{\lfloor \varepsilon n/(2k+\varepsilon)\rfloor}. \tag{5}$$

While this scaling is expected for systems with spatially local interactions, it is remarkable that this also holds for systems with all-to-all interactions: it implies that even by adding arbitrary all-to-all interactions to a spatially local Hamiltonian, one cannot raise the energy of too many states by a too large amount. One other important consequence of the bound (5) is that a quantum algorithm with a performance similar to that of Theorem 1 can then also be obtained by adapting known algorithms based on Grover search, by exploiting that the time Grover search needs to succeed decreases for a space which has more solutions.

Finally, on interaction graphs of maximum degree t, the bounds of Theorem 1, Theorem 2, and Equation (5) can in general be improved by replacing each occurence of k with $(1 - 1/t)k$.

Construction of H'—We start by constructing the new Hamiltonian H'. We will assume w.l.o.g. that all h_α act on exactly k qubits [29]. For each site $s = 1, \ldots, n$, define $\mathcal{I}(s) \coloneqq \{\alpha | h_\alpha \text{ acts on } s\}$ as the set of all interactions $\alpha \in \{1, \ldots, m\}$ which act on s. Further, define

$$e(s) \coloneqq \sum_{\alpha\in\mathcal{I}(s)} \|h_\alpha\|. \tag{6}$$

It quantifies the amount of energy $e(s)$ in H which originates in interactions which involve s. Importantly,

$$\sum_{s=1}^{n} e(s) = \sum_{\alpha=1}^{m} k\|h_\alpha\| = kM. \tag{7}$$

Since the labeling $s = 1, \ldots, n$ of sites is arbitrary, we choose to label them such that $e(s)$ is monotonously increasing, $e(1) \le e(2) \le \cdots \le e(n)$.

Now pick some $r \in \{1, \ldots, n\}$, and split the sites into the intervals $R \coloneqq \{1, \ldots, r\}$ and $\bar{R} \coloneqq \{r+1, \ldots, n\}$. The idea is to construct H' by only keeping interactions h_α which lie entirely within $\bar{R}$. Specifically, with T the set of all interactions α which only involve sites in $\bar{R}$, we define

$$H' \coloneqq \sum_{\alpha\in T} h_\alpha. \tag{8}$$

H' can either be regarded as a Hamiltonian acting on all n sites, or as acting only on the sites in $\bar{R}$, in which case we denote it by $H'_{\bar{R}}$, that is, $H' \equiv \mathbb{1}_R \otimes H'_{\bar{R}}$.

How well does H' approximate H? Let $|\phi\rangle$ be an arbitrary state. Then,

$$|\langle\phi|H - H'|\phi\rangle| \le \sum_{\alpha\notin T} |\langle\phi|h_\alpha|\phi\rangle| \le \sum_{\alpha\notin T} \|h_\alpha\|. \tag{9}$$

To bound the rightmost term, note first that

$$\sum_{\alpha\notin T} \|h_\alpha\| \le \sum_{s=1}^{r} e(s), \tag{10}$$

as the right hand side contains each $\|h_\alpha\|$ for $\alpha \notin T$ *at least* once. (On interaction graphs of maximum degree t, the right hand side of Eq. (10) can in general be improved by a multiplicative factor of $1 - 1/t$, which replaces k with $(1 - 1/t)k$ in all subsequent run-times stated in this paper; see Supplemental Material [30] for details.) Since $e(s)$ is monotonously increasing,

$$(n-r)e(r) \le \sum_{s=r+1}^{n} e(s) \le \sum_{s=1}^{n} e(s) \overset{(7)}{=} kM,$$

which yields $e(r) \le kM/(n - r)$. This allows us to bound

$$\sum_{s=1}^{r} e(s) \le re(r) \le \frac{r}{n-r} kM,$$

which together with (9) and (10) yields the bound

$$|\langle\phi|H - H'|\phi\rangle| \le \frac{r}{n-r} kM.$$

For a given $\delta > 0$ [31], we can thus obtain the bound

$$|\langle\phi|H - H'|\phi\rangle| \le \delta M \tag{11}$$

on how much H and H' can differ by choosing

$$r = \left\lfloor \frac{\delta n}{k+\delta} \right\rfloor. \tag{12}$$

From (11), we immediately obtain that the ground state energies $E_0(H)$ and $E_0(H')$ are close: To this end, let $|\Phi'\rangle$ be a ground state of H'. Since $\langle\Phi'|H|\Phi'\rangle - \langle\Phi'|H'|\Phi'\rangle \le \delta M$, we have

$$\langle \Phi'|H|\Phi'\rangle \leq \langle \Phi'|H'|\Phi'\rangle + \delta M = E_0(H') + \delta M. \quad (13)$$

Then, $E_0(H) \leq \langle \Phi'|H|\Phi'\rangle \leq E_0(H') + \delta M$, and correspondingly the other way around, which yields

$$|E_0(H) - E_0(H')| \leq \delta M. \quad (14)$$

Construction of low-energy states and density of states—As we have observed, we can consider H' as a Hamiltonian $H'_{\bar{R}}$ defined on the $n - r$ qubits in $\bar{R}$. Let $|\varphi\rangle_{\bar{R}}$ be a ground state of $H'_{\bar{R}}$. For an arbitrary state $|\chi\rangle_R$ on the r qubits in R, define $|\Phi''\rangle \coloneqq |\chi\rangle_R \otimes |\varphi\rangle_{\bar{R}}$. Then,

$$\langle \Phi''|H'|\Phi''\rangle = \langle \varphi|H'_{\bar{R}}|\varphi\rangle = E_0(H'). \quad (15)$$

That is, $|\Phi''\rangle$ is a ground state of H' (and as such satisfies (13), and thus

$$\langle \Phi''|H|\Phi''\rangle \overset{(13)}{\leq} E_0(H') + \delta M \overset{(14)}{\leq} E_0(H) + 2\delta M. \quad (16)$$

Note that given $|\varphi\rangle_{\bar{R}}$, $|\Phi''\rangle$ can be efficiently constructed, since all we have to do is to tensor it with an arbitrary state $|\chi\rangle_R$ defined on the qubits in R.

We just saw that there is a 2^r-dimensional space of states V—spanned by all $|\chi\rangle_R \otimes |\varphi\rangle_{\bar{R}}$ with arbitrary $|\chi\rangle_R$—with energy at most $2\delta M$ above $E_0(H)$. The min-max theorem (see, e.g., [32]) then implies that the 2^r smallest eigenvalues of H are all upper bounded by $E_0(H) + 2\delta M$. By replacing δ with $\varepsilon/2$ in (12), we get the claimed bound

$$\mathcal{C}(H, \varepsilon) \geq 2^r = 2^{\lfloor \varepsilon n/(2k+\varepsilon) \rfloor} \quad (17)$$

on the number of eigenstates with energy at most $E_0(H) + \varepsilon M$ [33].

Estimating the ground energy—The bound (17) can be used to improve the complexities of known ground energy estimation algorithms (e.g., the algorithm by Lin and Tong [27]), by exploiting that the run-time of quantum amplitude amplification with a target space of dimension $\mathcal{C}(H, \varepsilon)$ improves with its inverse square root [35]. However, we will give a simpler and more direct quantum algorithm, which moreover achieves (slightly) better performance.

The idea is simple: recall we wish to estimate $E_0(H)$ within error εM. We will use that $E_0(H') = E_0(H'_{\bar{R}})$ and apply the algorithm from [27] to $H'_{\bar{R}}$, which acts on $n - r$ qubits, to obtain an estimate $\hat{E}'$ of $E_0(H')$. This in turn gives a good estimate of $E_0(H)$ by (14). The computational speedup comes from working on a smaller system: the run-time of the algorithm now depends on $n - r$ rather than n.

We now work out the details of the algorithm. In (11) and (12), we choose $\delta = [1 - (1/n)]\varepsilon$. We then compute an estimate $\hat{E}'$ of $E_0(H'_{\bar{R}}) = E_0(H')$ such that

$$|\hat{E}' - E_0(H')| \leq \frac{\varepsilon}{n} M. \quad (18)$$

This can be done in

$$O^*\left(2^{(n-r)/2}\left(\frac{\varepsilon}{n}M\right)^{-1}\right) \overset{(12)}{=} O^*\left(2^{(1-\frac{\varepsilon}{k+\varepsilon})\frac{n}{2}}\varepsilon^{-1}\right) \quad (19)$$

time by applying the quantum algorithm from [27] using the completely mixed state (which has overlap $(\mathrm{Tr}\{[(I)/2^{n-r}]|\phi\rangle\langle\phi|\})^{1/2} = 2^{-(n-r)/2}$ with the ground state of $H'_{\bar{R}}$ and can easily be prepared) as initial state. Correspondingly, we can also compute $\hat{E}'$ on a classical computer using the Lanczos method [19,20] in time

$$O^*\left(2^{(n-r)}\left(\frac{\varepsilon}{n}\right)^{-1}\right) \overset{(12)}{=} O^*\left(2^{(1-\frac{\varepsilon}{k+\varepsilon})n}\varepsilon^{-1}\right). \quad (20)$$

The returned estimate will with high probability satisfy (18), which together with (14) [recall that $\delta = [1 - (1/n)]\varepsilon$] guarantees that $|\hat{E}' - E_0(H)| \leq \varepsilon M$ holds, as desired.

Note that in order to apply the algorithm from [27] we need to construct a block-encoding of H'. Since H' is k-local, and thus sparse, a block-encoding can be implemented efficiently using the techniques from [36].

Preparing a low-energy state—We now discuss how to prepare a low-energy state of H. We will achieve this by first preparing a low-energy state of the Hamiltonian $H'_{\bar{R}}$ acting on $n - r$ qubits. As we have seen, $H'_{\bar{R}}$ trivially embeds into the Hamiltonian on the full n qubits, $H' \equiv \mathbb{1}_R \otimes H'_{\bar{R}}$. This embedding thus yields a low-energy state for H', and thus by (11) also for H.

To start, choose $\delta = [1 - (1/n)](\varepsilon/2)$ in (11) and (12). Let

$$H'_{\bar{R}} = \sum_{i=1}^{2^{n-r}} \lambda_i |\phi_i\rangle\langle\phi_i|$$

be the eigenvalue decomposition of $H'_{\bar{R}}$. Using the quantum singular value transformation (QSVT) [36,37], we can implement in $O^*(\log(1/\chi)\varepsilon^{-1})$ time an operator

$$\Pi = \sum_{i=1}^{2^{n-r}} P(\lambda_i)|\phi_i\rangle\langle\phi_i|,$$

where $P\colon [-M, M] \to [0, 1]$ is a polynomial such that

$$P(\lambda_i) \geq 1 - \chi \quad \text{if } \lambda_i \leq E_0(H'), \quad (21)$$

$$P(\lambda_i) \leq \chi \quad \text{if } \lambda_i \geq E_0(H') + \frac{\varepsilon}{n}M. \quad (22)$$

The operator Π is (to high accuracy) a projection onto the low-energy space of $H'_{\bar{R}}$. For simplicity, in the analysis below we will assume $\chi = 0$. The QSVT framework from [36,37] does not directly implement Π. Instead, it implements a block-encoding of Π, i.e., a larger unitary matrix U such that Π appears (possibly after normalization) as the top left block of U:

$$U = \begin{pmatrix} \Pi & \cdot \\ \cdot & \cdot \end{pmatrix}.$$

The key property of this block-encoding is as follows: For any state σ on $n - r$ qubits, applying U on $\sigma \otimes (|0\rangle\langle 0|)^{\otimes \ell}$, where ℓ denotes the number of ancillas used by U, and subsequently measuring the ancillas in the computational basis, returns—with probability $\mathrm{Tr}[\Pi\sigma]$—the state

$$\frac{\Pi\sigma\Pi}{\mathrm{Tr}[\Pi\sigma]} \tag{23}$$

on the first $n - r$ qubits; i.e., it successfully applies Π on σ with probability $\mathrm{Tr}[\Pi\sigma]$. Due to (14) and (22), the state of (23) has low energy:

$$\mathrm{Tr}\left[H'_{\bar{R}}\left(\frac{\Pi\sigma\Pi}{\mathrm{Tr}[\Pi\sigma]}\right)\right] \leq E_0(H) + \left(\delta + \frac{\varepsilon}{n}\right)M. \tag{24}$$

In order to prepare a low-energy state of H, we take σ as the completely mixed state σ_0 on $n - r$ qubits, i.e. $\sigma_0 = (1/2^{n-r})\sum|\phi_i\rangle\langle\phi_i|$; note that this state can be prepared efficiently. As above, we apply the unitary U and measure the ancilla qubits. The probability that we obtain the state (23) is

$$\mathrm{Tr}[\Pi\sigma_0] = \sum_{i=1}^{2^{n-r}} \frac{\mathrm{Tr}(\Pi|\phi_i\rangle\langle\phi_i|)}{2^{n-r}} \geq \frac{1}{2^{n-r}}, \tag{25}$$

where the inequality follows from (21), since there exists at least one $|\phi_i\rangle$ for which $\lambda_i = E_0(H')$. Finally, we embed this state in the full n qubits by adding r qubits each initialized to $|0\rangle\langle 0|$:

$$\rho \coloneqq \left((|0\rangle\langle 0|)^{\otimes r}\right)_R \otimes \left(\frac{\Pi\sigma_0\Pi}{\mathrm{Tr}[\Pi\sigma_0]}\right)_{\bar{R}}. \tag{26}$$

This state satisfies Theorem 2,

$$\mathrm{Tr}[H\rho] \overset{(11)}{\leq} \mathrm{Tr}[H'\rho] + \delta M \overset{(24)}{\leq} E_0(H) + \varepsilon M \tag{27}$$

[recall that $\delta = [1 - (1/n)](\varepsilon/2)$].

Using quantum amplitude amplification [35], the probability (25) can be amplified to a probability arbitrarily close to 1 using $O(2^{(n-r)/2})$ calls to U. Since the unitary U can be implemented in $O^*(\epsilon^{-1})$ time, the overall time complexity is

$$O^*\left(2^{(n-r)/2}\epsilon^{-1}\right) \overset{(12)}{=} O^*\left(2^{(1-\frac{\varepsilon}{2k+\varepsilon})\frac{n}{2}}\epsilon^{-1}\right). \tag{28}$$

Concluding remarks—In this Letter we gave a quantum algorithm which provides a superquadratic speedup for computing properties of a general Hamiltonian with all-to-all interactions and k-body terms. The main contribution of our work is arguably not the magnitude of run-time improvement itself, but rather the following two points: first, and perhaps most surprisingly, that the standard square root speedup attainable by unstructured Grover search can be beaten, without assumptions on the local Hamiltonian itself; and second, that our approach is not heuristic, but yields rigorous and worst-case run-time guarantees. In both regards, our result can be interpreted as a quantum version of the (classical) breakthrough result by Hirsch [38], which broke the natural bound based on brute-force search for hard classical optimization problems.

Hirsch's algorithm selects a random assignment and tries to improve it by repeatedly flipping a value of a variable chosen randomly from an unsatisfied clause (again chosen randomly). The classical analogue of our approach for computing approximately optimal solutions is similar, and slightly simpler than Hirsch: we repeatedly select a random assignment—the bound on the running time then follows from our bound on the dimension of the low energy space. As further discussed in Supplemental Material [30], the complexity of this approach is similar to the complexity of Hirsch's algorithm. As Hirsch's work has been influential to the development of rigorous classical approximation algorithms [39–43], we expect that our discoveries will initiate research on approximation algorithms for properties of general Hamiltonians.

In addition, the gap exponential time hypothesis (GAP-ETH) [44,45] states that there exist constants c, ε such that approximating ground state energies for 3-local Hamiltonians within error $\pm\epsilon M$ is classically impossible in $O(2^{cn})$ time. It is thus reasonable to posit a *quantum* GAP-ETH, which is identical except the no-go is for quantum algorithms in $O(2^{cn/2})$ time. In this sense, one would not expect a major improvement over our Theorem 1.

An open question is whether our approach can be used to produce a guiding state for the ground state, to be input into the algorithms of [26,27]. On the one hand, if one believes the quantum strong exponential time hypothesis (QSETH) [46], which states that Boolean satisfiability problems cannot be solved with superquadratic speedup over brute force search for large k, even by a quantum computer, then producing such a guiding state with superquadratic speedup is impossible. On the other hand, our run-time is best for small k, so it is not clear whether QSETH should pose an obstruction here. With this said, however, we show in Supplemental Material [30] that for any $K \geq \epsilon M$, the low-energy state we produce in (26) has $1/K$ overlap with the space of energy at most $E_0(H) + K$. By choosing $K = (1 + 1/p(n))\varepsilon M$ for a sufficiently large polynomial p, one effectively recovers the energy bound of Eq. (27), while maintaining that ρ has inverse polynomial overlap with the low energy space.

Acknowledgments—We thank Toby Cubitt for helpful discussions. Part of this work was done when visiting the Simons Institute for the Theory of Computing. S. G. is supported by DFG Grants No. 432788384 and No. 450041824, BMBF project PhoQuant (13N16103), and project PhoQC from the State of Northrhine Westphalia. Z. L. is supported by DOE QSA Grant No. FP00010905, NSF QLCI Grant No. 2016245, DOE grant No. DE-SC0024124 and NSF Grant No. 2311733. F. L. G. is supported by JSPS KAKENHI Grants No. JP20H05966 and No. 24H00071 and MEXT Q-LEAP Grant No. JPMXS0120319794. N. S. is supported by the Austrian Science Fund (FWF) [Grants DOI: 10.55776/COE1, DOI: 10.55776/P36305, and DOI: 10.55776/F71], the European Union–NextGenerationEU, and the European Union's Horizon 2020 research and innovation program through Grant No. 863476 (ERC-CoG SEQUAM). S. T. is supported by JSPS KAKENHI Grants No. JP20H05961, No. JP20H05967, and No. JP22K11909.

[1] A. Szabo and N. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*, Dover Books on Chemistry (Dover Publications, New York, 1996).

[2] T. Helgaker, P. Jorgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (Wiley, New York, 2014).

[3] M. A. Ruderman and C. Kittel, Indirect exchange coupling of nuclear magnetic moments by conduction electrons, Phys. Rev. **96**, 99 (1954).

[4] M. Saffman, T. G. Walker, and K. Mølmer, Quantum information with Rydberg atoms, Rev. Mod. Phys. **82**, 2313 (2010).

[5] R. Islam, C. Senko, W. Campbell, S. Korenblit, J. Smith, A. Lee, E. Edwards, C.-C. Wang, J. Freericks, and C. Monroe, Emergence and frustration of magnetic order with variable-range interactions in a trapped ion quantum simulator, Science **340**, 583 (2013).

[6] R. Blatt and C. F. Roos, Quantum simulations with trapped ions, Nat. Phys. **8**, 277 (2012).

[7] B. Yan, S. A. Moses, B. Gadway, J. P. Covey, K. R. Hazzard, A. M. Rey, D. S. Jin, and J. Ye, Observation of dipolar spin-exchange interactions with lattice-confined polar molecules, Nature (London) **501**, 521 (2013).

[8] D. Otten, S. Rubbert, J. Ulrich, and F. Hassler, Universal power-law decay of electron-electron interactions due to nonlinear screening in a Josephson junction array, Phys. Rev. B **94**, 115403 (2016).

[9] S. Gopalakrishnan, B. L. Lev, and P. M. Goldbart, Frustration and glassiness in spin models with cavity-mediated interactions, Phys. Rev. Lett. **107**, 277201 (2011).

[10] A. Kitaev, A. Shen, and M. Vyalyi, *Classical and Quantum Computation*, Graduate Studies in Mathematics (American Mathematical Society, Providence, 2002), Vol. 47.

[11] J. Kempe, A. Y. Kitaev, and O. Regev, The complexity of the local Hamiltonian problem, SIAM J. Comput. **35**, 1070 (2006).

[12] N. Bansal, S. Bravyi, and B. M. Terhal, Classical approximation schemes for the ground-state energy of quantum and classical Ising spin Hamiltonians on planar graphs, Quantum Inf. Comput. **9**, 701 (2009).

[13] D. Aharonov, I. Arad, and T. Vidick, Guest column: The quantum PCP conjecture, ACM SIGACT News **44**, 47 (2013).

[14] R. Oliveira and B. M. Terhal, The complexity of quantum spin systems on a two-dimensional square lattice, Quantum Inf. Comput. **8**, 0900 (2008).

[15] T. Cubitt and A. Montanaro, Complexity classification of local Hamiltonian problems, SIAM J. Comput. **45**, 268 (2016).

[16] S. Piddock and A. Montanaro, The complexity of antiferromagnetic interactions and 2D lattices, Quantum Inf. Comput. **17**, 636 (2017).

[17] S. Arora and S. Safra, Probabilistic checking of proofs: A new characterization of NP, ACM J. Emerging Technol. Comput. Syst. **45**, 70 (1998).

[18] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy, Proof verification and the hardness of approximation problems, ACM J. Emerging Technol. Comput. Syst. **45**, 501 (1998).

[19] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, J. Res. Natl. Bur. Stand. **45**, 255 (1950).

[20] J. Kuczyński and H. Woźniakowski, Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start, SIAM J. Matrix Anal. Appl. **13**, 1094 (1992).

[21] R. Impagliazzo and R. Paturi, On the complexity of k-SAT, J. Comput. Syst. Sci. **62**, 367 (2001).

[22] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the 28th Annual ACM Symposium on Theory of Computing* (1996), pp. 212–219.

[23] D. Poulin and P. Wocjan, Preparing ground states of quantum many-body systems on a quantum computer, Phys. Rev. Lett. **102**, 130503 (2009).

[24] J. van Apeldoorn, A. Gilyén, S. Gribling, and R. de Wolf, Quantum SDP-solvers: Better upper and lower bounds, Quantum **4**, 230 (2020).

[25] A. Kerzner, V. Gheorghiu, M. Mosca, T. Guilbaud, F. Carminati, F. Fracas, and L. Dellantonio, A square-root speedup for finding the smallest eigenvalue, arXiv:2311.04379.

[26] Y. Ge, J. Tura, and J. I. Cirac, Faster ground state preparation and high-precision ground energy estimation with fewer qubits, J. Math. Phys. (N.Y.) **60**, 022202 (2019).

[27] L. Lin and Y. Tong, Near-optimal ground state preparation, Quantum **4**, 372 (2020).

[28] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu, M. Kastoryano, R. Babbush, J. Preskill, D. R. Reichman, E. T. Campbell,

E. F. Valeev, L. Lin, and G. K.-L. Chan, Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry, Nat. Commun. **14**, 1952 (2023).

[29] Interactions acting on less qubits can be "padded" by adding further qubits on which they act trivially.

[30] See Supplemental Material at http://link.aps.org/supplemental/10.1103/29qw-bssx for improved runtime analyses, guiding state analyses, and relationship with Hirsch's work.

[31] While any value of $\delta > 0$ can be chosen, when δ is very small we obtain $r = 0$, which does not give any speedup (since we are left with the original Hamiltonian).

[32] R. A. Horn and C. R. Johnson, *Matrix Analysis* (Cambridge University Press, Cambridge, England, 1985).

[33] Note that Cubitt and González-Guillén [34] previously gave exponential lower bounds on the low-energy spaces of local Hamiltonians, but only for *history state* Hamiltonians (such as Kitaev's circuit-to-Hamiltonian construction [10]). Our result, in contrast, applies to any k-local Hamiltonian.

[34] C. E. González-Guillén and T. S. Cubitt, History-state Hamiltonians are critical, arxiv:1810.06528.

[35] G. Brassard, P. Høyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, Contemp. Math. **305**, 53 (2002).

[36] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (2019), pp. 193–204.

[37] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang, Grand unification of quantum algorithms, PRX Quantum **2**, 040203 (2021).

[38] E. A. Hirsch, Worst-case study of local search for MAX-k-SAT, Discrete Appl. Math. **130**, 173 (2003).

[39] B. Escoffier, V. T. Paschos, and E. Tourniaire, Approximating MAX SAT by moderately exponential and parameterized algorithms, Theor. Comput. Sci. **560**, 147 (2014).

[40] J. Alman, T. M. Chan, and R. R. Williams, Faster deterministic and las vegas algorithms for offline approximate nearest neighbors in high dimensions, in *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms* (2020), pp. 637–649.

[41] A. Drucker, An improved exponential-time approximation algorithm for fully-alternating games against nature, in *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science* (2020), pp. 1081–1090.

[42] T. Korhonen, A single-exponential time 2-approximation algorithm for treewidth, in *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science* (2021), pp. 184–192.

[43] B. C. Esmer, A. Kulik, D. Marx, D. Neuen, and R. Sharma, Optimally repurposing existing algorithms to obtain exponential-time approximations, in *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms* (2024), pp. 314–345.

[44] I. Dinur, Mildly exponential reduction from gap-3SAT to polynomial-gap label-cover, Electronic colloquium on computational complexity ECCC; research reports, surveys and books in computational complexity (2016).

[45] P. Manurangsi and P. Raghavendra, A birthday repetition theorem and complexity of approximating dense CSPs, in *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 80, edited by I. Chatzigiannakis, P. Indyk, F. Kuhn, and A. Muscholl (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2017), pp. 78:1–78:15.

[46] H. Buhrman, S. Patro, and F. Speelman, A framework of quantum strong exponential-time hypotheses, in *Proceedings of the 38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 187, edited by M. Bläser and B. Monmege (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2021), pp. 19:1–19:19.