ARTICLES
PUBLISHED ONLINE: 1 JUNE 2015 | DOI: 10.1038/NPHYS3345
nature physics

A polynomial time algorithm for the ground state of one-dimensional gapped local Hamiltonians

Zeph Landau[1], Umesh Vazirani[1] and Thomas Vidick[2]*

The density matrix renormalization group method has been extensively used to study the ground state of 1D many-body systems since its introduction two decades ago. In spite of its wide use, this heuristic method is known to fail in certain cases and no certifiably correct implementation is known, leaving researchers faced with an ever-growing toolbox of heuristics, none of which is guaranteed to succeed. Here we develop a polynomial time algorithm that provably finds the ground state of any 1D quantum system described by a gapped local Hamiltonian with constant ground-state energy. The algorithm is based on a framework that combines recently discovered structural features of gapped 1D systems with an efficient construction of a class of operators called approximate ground-state projections (AGSPs). The combination of these tools yields a method that is guaranteed to succeed in all 1D gapped systems. An AGSP-centric approach may help guide the search for algorithms for more general quantum systems, including for the central challenge of 2D systems, where even heuristic methods have had more limited success.

A fundamental obstacle for the study of quantum many-body systems, such as those involved in high-temperature superconductivity[1] or the fractional quantum Hall effect[2,3], is the sheer number of parameters required to give a complete description of their quantum state. The dimension of the Hilbert space required to accommodate all such states increases exponentially with the number of particles, and systems with more than a few dozen particles already present an insurmountable challenge for numerical simulation through exact diagonalization. A simple counting argument reveals that states that can be obtained as the equilibrium state of a nearest-neighbour Hamiltonian on a low-dimensional lattice can occupy only a tiny (polynomial) fraction of this exponential space, and the resulting manifold has been called the 'physically relevant corner of Hilbert space'[4]. This observation by itself is of limited practical use as it does not provide a workable description of the relevant quantum states. An important program thus remains to be carried out to provide an effective map of the physical corner of Hilbert space. First, devise compact methods to represent the relevant states that also allow efficient computation of expectation values of local observables. Second, show that the compact description can itself be efficiently computed from the Hamiltonian.

In this work we address the second part of this challenge: we provide the first algorithmic method that is both efficient and provably correct for finding compact representations of a large class of many-body systems. The first part of the challenge has been very successfully addressed in recent years by the use of tensor networks and their low-dimensional specializations such as matrix-product states[5] (MPSs; see Fig. 1) in one dimension, and projected-entangled pair states[6] (PEPSs) in two dimensions. These impressive successes are directly related to the phenomenon of entanglement; indeed, a tensor network representation necessarily implies that the associated quantum state has small entanglement rank. Justification for this assumption comes from a sweeping conjecture in condensed-matter physics, called the area law, which asserts that gapped Hamiltonians have limited entanglement in their ground states[7]. More precisely, it asserts that for any subset S of particles, the entanglement entropy between S and $\bar{S}$ is bounded by the surface area of S rather than the trivial bound of the volume of S. A seminal result of Hastings[8] rigorously established the area law for one-dimensional (1D) gapped Hamiltonians. As a consequence he proved that the ground state of such systems is accurately described by polynomial-size MPSs. Finite-temperature Gibbs states in any dimension are also known to satisfy an area law[9], and have recently been shown to admit a polynomial-size PEPS representation[10] (which implies, but is not implied by ref. 11, an area law for the mutual information). The area law for ground states of gapped Hamiltonians in more than one dimension remains a major open problem.

The second part of the challenge has seen very few rigorous results. There is essentially a single method used to search for MPS representations, the density matrix renormalization group[12] (DMRG), which has been very successful in practice in its original setting of 1D systems at zero temperature[13,14]. Much effort has been devoted to extending the method to, for example, systems in two dimensions[15] or systems out of equilibrium[16], but many challenges remain[17,18]. In some cases, including certain instances of 1D gapped systems, however, one will find that the method gets stuck in local optimum. It has also been demonstrated that a widely used variant of DMRG, which implements a two-site optimization procedure, must in the worst case solve NP-hard computational problems[19] to completely solve the local optimization problem. This implies that for such variants efficiency inevitably comes at the cost of performing ungrounded approximations. To avoid these pitfalls DMRG researchers have developed an increasing range of practical heuristics[20]. Unfortunately, making the right choice of parameters (such as discretization error, number of sweeps, bond dimension, initial state) and heuristic inevitably requires some a priori knowledge, or at least intuition, on the physical system—a task that DMRG is arguably designed to solve in the first place. Moreover, in the absence of guarantees one can never be sure that the solution returned is a good approximation to the optimum.

[1]Computer Science Division, University of California, Berkeley, Berkeley, California 94720, USA. [2]Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, California 91125, USA. *e-mail: vidick@cms.caltech.edu

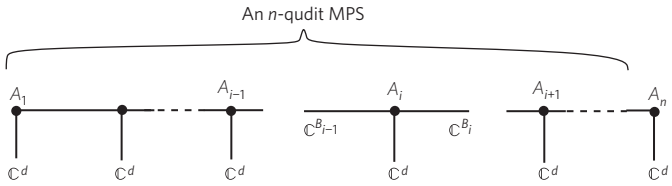

Figure 1 | Entanglement and MPSs. A bipartite state $|v\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ has a unique decomposition $|v\rangle = \sum_{j=1}^{B} \lambda_j |a_j\rangle|b_j\rangle$, with $\{|a_j\rangle\}$ (respectively $\{|b_j\rangle\}$) being a family of orthonormal vectors of $\mathcal{H}_A$ (respectively $\mathcal{H}_B$). The entanglement rank (also called the Schmidt rank) of $|v\rangle$ is the number of terms, B, in the decomposition, and the entanglement entropy is $S = \sum_j \lambda_j^2 \log(1/\lambda_j^2)$. An MPS representation of a state of n qudits is specified by a sequence of tensors $A_1, \ldots, A_n$, where $A_1 \in \mathbb{C}^d \times \mathbb{C}^{B_1}$, $A_i \in \mathbb{C}^{B_{i-1}} \times \mathbb{C}^d \times \mathbb{C}^{B_i}$ for $1 \leq i < n$, and $A_n \in \mathbb{C}^{B_{n-1}} \otimes \mathbb{C}^d$. The bond dimension B_i corresponds to the entanglement rank across the $(i, i+1)$ cut.

Our work addresses these drawbacks by providing an efficient algorithm that in all cases is guaranteed to return an MPS representation for a very good approximation to the ground state of any gapped 1D system. Our algorithm makes a single pass over the chain and provably constructs a polynomial-size representation for a very close (inverse-polynomial) approximation to the ground state; and it is guaranteed to never get stuck. Letting N be the length of the chain, ε the gap of the Hamiltonian and δ the desired accuracy, the algorithm runs in time $N^{c(\epsilon)}\mathrm{poly}(N/\delta)$, where $c(\epsilon) = 2^{O(1/\epsilon)}$ is an exponent that increases with the gap but is constant for constant spectral gap. (See Methods for a description of the precise dependence.)

To avoid the local minima that have plagued DMRG, we introduce a new procedure that improves any approximation to the ground state by acting with an operator, called approximate ground-state projection (AGSP), that is derived from the Hamiltonian. AGSPs were introduced recently in the context of analysing the ground-state entanglement of 1D systems[21]. For the algorithm, a new type of AGSP, called a sampling AGSP, is constructed out of the local Hamiltonian describing the quantum system. The AGSP has no analogue in DMRG and is the major new ingredient that allows our algorithm to systematically avoid the pitfalls faced by DMRG, including the latter's difficulties in avoiding local minima.

To better appreciate the approach taken by our algorithm, it is useful to draw an analogy. Think of the Hilbert space as a vast lake that is lightly frozen—the rare solidly frozen stretches corresponding to those states that have a succinct MPS description. The goal of the algorithm is to pull a sled to a distant island, representing the actual ground state, during a blizzard. Success requires navigating a route that at all times keeps the sled on solidly frozen ice. The low visibility forces the algorithm to map its route one step at a time. At any iteration, the algorithm restricts its exploration to a polynomial dimensional subspace of the Hilbert space, represented by a set of MPSs describing a basis for the space. It remains on solid ice by keeping the size of these MPS representations bounded through bond trimming. The dimension of the Hilbert space that must be explored in each iteration is kept in check by appealing to a crucial property of 1D systems called decoupling. The trickiest part is navigation. Unavoidable errors in navigation at each step can add up to throw the algorithm totally off course in the exponentially large Hilbert space. The AGSP plays the central role of navigator, guiding the search and keeping the navigation errors under control.

Hamiltonian complexity

Inspiration for our results comes from a productive line of work that has brought the perspective of complexity theory to bear on problems inspired by condensed-matter physics. The area was initiated by Kitaev's seminal result[22] showing that the problem of computing ground states of local Hamiltonians is, in general, QMA-hard. QMA is the quantum analogue of the class NP, and it is believed to contain problems that are strictly harder than NP-complete problems such as 3SAT. A proof of QMA-hardness provides strong evidence not only that the classical description of ground states of local Hamiltonians must be hard to find, but also that no useful small-size description even exists. Further extensions to 2D lattices[23], 1D Hamiltonians[24] and even translation-invariant systems[25] show that in the absence of any further condition on the Hamiltonian none of these systems is amenable to a systematic analysis.

These complexity-theoretic results directly clash with the practical successes described above and suggest a search for the structural features that delineate the separation between Hamiltonians whose ground states can efficiently be found, and those for which the problem is hard. Although it is tempting to conjecture that the existence of an area law should be a sufficient condition, this possibility is ruled out by examples of 1D systems in which the ground state satisfies an area law with only a logarithmic correction but for which computing any approximating MPS is as hard as integer factoring[26], which is widely believed to require exponential time: the computational intractability of integer factoring is a standard assumption underlying much of modern cryptography including the RSA crypto system.

Recent work by Arad *et al.*[27] showed that the gap condition rules out such examples. They gave an exponential improvement over Hastings' area law bound, and in the process also established that ground states of gapped 1D Hamiltonians have additional structure: they can be well approximated by an MPS whose bond dimension has a sub-linear dependence on n. Combined with the exponential-time dynamic programming-based algorithm[28,29] this led to a sub-exponential algorithm for gapped 1D systems. In this work we demonstrate that the gap condition can actually be used to construct a provably efficient (polynomial time) algorithm.

The algorithm

The problem of finding an (approximation to the) ground state can be expressed as a convex optimization problem: $\min \mathrm{tr}(H\sigma)$ subject to $\mathrm{tr}(\sigma) = 1$, $\sigma \geq 0$, where σ is an n-qubit density matrix describing the state. As the dimension of σ scales exponentially in n, solving this convex program generically requires exponential time. The driving idea behind our algorithm is to progressively construct, in a single left→right pass, a basis for a subspace of polynomial dimension of the whole space guaranteed to contain a suitable approximation to the ground state $|\Gamma\rangle$. This contrasts strongly with the usual approach, taken by DMRG, of performing multiple passes of local optimization with no guarantee of the quality of the approximation reached after any number of passes.

The subspace constructed at the ith step is specified by a small spanning set of states on the first i qudits, each specified by an MPS of fixed polynomial size. Such a set S is called an (i, δ)-viable set for $|\Gamma\rangle$ if there exists a state $|\phi\rangle \in (\mathbb{C}^d)^{\otimes n}$ such that $|\langle\phi|\Gamma\rangle| \geq 1 - \delta$ and such that the reduced density of $|\phi\rangle$ on the first i qudits is supported on $\mathrm{Span}(S)$. Here, $|\phi\rangle$ is called a witness for S, and δ the error of S. Two other important parameters associated with a viable set S are its cardinality s, and the maximum bond dimension B for the MPS representations of the elements of S.

The algorithm iteratively constructs a sequence of polynomial-size (i, δ)-viable sets $S_0, \ldots, S_i, \ldots, S_n$, for suitably small δ, from left to right, extending by one qudit in each iteration. The construction of $S_0 = \{1\}$ is trivial. Moreover, given S_n the ground state can be efficiently found by solving the convex program outlined above, where the optimization is over all σ supported on the span of S_n. It remains to show how S_{i-1} can be extended to S_i while maintaining reasonable bounds on the parameters of these viable sets. This is

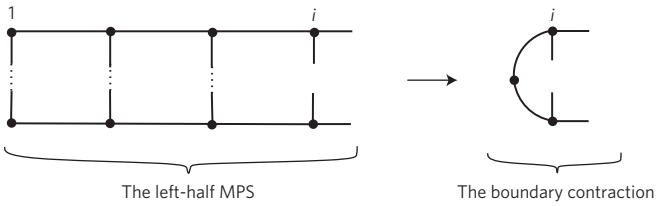

Figure 2 | Boundary contraction associated with a left-half MPS on the first *i* qudits with an open bond.

performed in a sequence of four steps: extension, size and bond trimming, and error reduction.

The first step, extension, is simple: it replaces S_{i-1} by the set $S_i^{(1)}$ obtained by extending states in S_{i-1} by taking their tensor product with a basis of states for the space of the ith qudit: $S_i^{(1)} := \{|s\rangle \otimes |j\rangle : |s\rangle \in S_{i-1}, 1 \leq j \leq d\}$. This step results in a multiplicative factor d increase in the cardinality of the viable set. Although this may not seem like a large increase, left unchecked it would result in an exponential factor d^n over the n iterations. In the blizzard analogy, this increase in cardinality corresponds to the accumulating snow that progressively weighs down the sled. The next step corresponds to shovelling off some of the accumulated snow from the sled.

The goal of the second step, size trimming, is to reduce the cardinality of $S_i^{(1)}$ to create a smaller set $S_i^{(2)}$. To achieve this we rely on the notion of boundary contraction of a state, which captures information representing how the components of the state associated with the left and right halves of the chain are combined together (see Fig. 2). Subject to a fixed boundary contraction, the problem of finding a state of minimal energy can be decoupled into two disjoint problems—to the left and to the right of the cut.

If we knew the correct boundary contraction, associated with the ground state $|\Gamma\rangle$, it would suffice to include in $S_i^{(2)}$ any state in $S_i^{(1)}$ that has that boundary contraction and has minimal energy among all such states. As the correct boundary contraction is unknown we need to include the above for every possible boundary contraction, taken from a suitably discretized ϵ-net. This approach, already taken in the exponential-time algorithm[28,29], encounters a major difficulty: the size of the ϵ-net is necessarily exponential in the dimension of the bond across sites i and $i+1$, and the 1D area law[8] only guarantees a polynomial bond dimension for inverse-polynomial approximations to $|\Gamma\rangle$.

To overcome this difficulty we appeal to a stronger structural property that follows from existing proofs of the 1D area law[8,27]: for any given cut and constant η there exists a state having a constant bond dimension

$$B_\eta = \exp(O((1/\eta)\log^3 d \log(1/\eta)))$$

across that cut (and polynomial across all others) that is an η-approximation to the ground state $|\Gamma\rangle$. This implies that, for the purpose of obtaining a constant approximation, it suffices to consider a polynomial-size net over the space of boundary contractions with constant bond dimension B. Letting $\mathcal{N}$ be such a net (we give an explicit construction in the Supplementary Information), for each $X \in \mathcal{N}$, we find a state supported on $\mathrm{Span}\{S_i^{(1)}\} \otimes \mathbb{C}^B$ of minimum energy among those states whose boundary contraction is close to X. This can be done efficiently by solving an appropriate convex program (see Methods for a complete description of the procedure).

As a result, both the bond dimension B of MPS representations in $S_i^{(2)}$ and the approximation error δ have increased: with our choice of parameters δ is now a constant, $\delta = 1/12$. In the blizzard analogy, this respectively corresponds to the sled now being on thinner ice (increased B) and it being pushed off course (bigger δ). The goal of the remaining two steps is to reset these parameters.

In the third step, bond trimming, the algorithm trims each of the bonds $1, \ldots, i-1$ of every $|u\rangle \in S_i^{(2)}$ to some fixed polynomial by iteratively keeping only the components associated with the largest Schmidt coefficients across each bond. Keeping a sufficiently large polynomial number of coefficients does not affect the approximation error δ by too much. The resulting set $S_i^{(3)}$ is a $(i, 1/2)$-viable set whose cardinality and maximal bond dimension are both bounded by fixed polynomials. With the sled back on solid ice, we turn to the most challenging task: controlling navigation error.

At this stage, the viable set $S_i^{(3)}$ contains the left Schmidt vectors of a δ-approximation $|\psi\rangle$ to the ground state $|\Gamma\rangle$, with $\delta = 1/2$. To complete the fourth step it remains to improve the approximation to inverse polynomial. For this the key idea is to use a well-designed operator K, the AGSP. This AGSP is the main workhorse of the algorithm. It has two important properties.

First, for any vector $|\psi\rangle$ such that $|\langle\Gamma|\psi\rangle| \geq 1/2$ it holds that $|\langle\Gamma|K|\psi\rangle| \geq 1 - 1/\mathrm{poly}(n)$. Second, the application of K to any MPS with bond dimension B can be performed efficiently and results in an MPS with bond dimension at most $p(n)B$, where $p(n)$ is a fixed polynomial, depending only on the inverse gap ε^{-1} of H and the local dimension d. We give a new efficient construction of an AGSP, the sampling AGSP, that satisfies both conditions. The construction is described in the Methods; it is based on a randomized sampling procedure applied to the term-by-term expansion of the (appropriately scaled) operator $(\mathrm{Id} - H/n)^m$. We will apply the AGSP to each state in $S_i^{(3)}$. Doing so requires some care, as a priori K acts on the whole chain, but states in $S_i^{(r)}$ are only defined on the first i qudits. For this we decompose K across the $(i, i+1)$ cut as a sum $K = \sum A_j \otimes B_j$ over polynomially many terms, such that furthermore each A_j, which acts on the first i qudits, has polynomial Schmidt rank across every cut. Using this decomposition the set $S_i^{(4)} := \{A_j|s\rangle : |s\rangle \in S_i^{(3)}\}$ can be efficiently computed and shown to be a $(i, \delta = 1/\mathrm{poly}(n))$ viable set. This completes the fourth and last step of the ith iteration of the algorithm.

Concluding remarks

The running time of our algorithm, although polynomial in the system size, is hardly practical as the exponent of this polynomial given by our analysis is very large. An important path for further research consists of improving this exponent and devising a practically efficient method that provides the same theoretical guarantees. In this respect it would be worthwhile to explore whether some of the new components introduced in our algorithm—such as the use of the sampling AGSP or the notion of viable sets—could be folded into existing heuristic algorithms, such as DMRG, to improve their success rate on difficult instances. For instance, the use of the sampling AGSP suggests a direct method to improve the quality and convergence speed of DMRG by randomly applying terms from the Hamiltonian to the solutions constructed at each step. At a high level this amounts to applying an efficient, approximate variant of time evolution, for short times, to the partial solutions, and suggests an efficient composite between DMRG and the iTEBD algorithm[30].

The most promising direction for further research is the application of the ideas introduced here to the study of higher-dimensional systems. In particular the use of the sampling AGSP as a major driving force in our algorithm coupled with the recent use of different AGSPs to prove an improved area law for 1D systems[27] firmly establishes the AGSP as a powerful new tool. This suggests a new approach, with AGSPs playing the central role, to obtaining algorithms for 2D systems, where despite great recent progress[4,6,31–39] the methods have enjoyed limited success as compared with their 1D counterparts.

In more than one dimension the formalism of MPSs has been fruitfully generalized to PEPSs (refs 40,41), which are conjectured

to capture the physics of quantum gapped phases[42]. States that have PEPS representations satisfy an area law by definition, and a parent Hamiltonian can be associated with any PEPS (ref. 43). This puts us in the same situation as in one dimension, under the a priori assumption that an area law holds. A substantial difference between PEPSs and MPSs, however, is that the former do not seem to allow the efficient computation of local observables. This opens up a number of challenges for which the algorithmic ideas introduced in this work may prove helpful—for example, the 2D situation can naturally be mapped to a 1D problem by applying the transfer matrix formalism to represent a 2D network as the result of successive compositions, along the vertical axis, of a horizontal 1D matrix-product operator[44].

Methods

Methods and any associated references are available in the online version of the paper.

Received 28 October 2014; accepted 23 April 2015; published online 1 June 2015

References

1. Bednorz, J. G. & Müller, K. A. Possible high-t_c superconductivity in the Ba–La–Cu–O system. *Z. Phys. B* **64,** 189–193 (1986).
2. Tsui, D. C., Stormer, H. L. & Gossard, A. C. Two-dimensional magnetotransport in the extreme quantum limit. *Phys. Rev. Lett.* **48,** 1559–1562 (1982).
3. Laughlin, R. B. Anomalous quantum Hall effect: An incompressible quantum fluid with fractionally charged excitations. *Phys. Rev. Lett.* **50,** 1395–1398 (1983).
4. Verstraete, F. & Cirac, J. I. Renormalization algorithms for quantum-many body systems in two and higher dimensions. Preprint at http://arXiv.org/abs/cond-mat/0407066 (2004).
5. Vidal, G. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.* **91,** 147902 (2003).
6. Verstraete, F., Murg, V. & Cirac, J. I. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Adv. Phys.* **57,** 143–224 (2008).
7. Eisert, J., Cramer, M. & Plenio, M. B. Colloquium. *Rev. Mod. Phys.* **82,** 277–306 (2010).
8. Hastings, M. B. An area law for one-dimensional quantum systems. *J. Stat. Mech.* **2007,** P08024 (2007).
9. Wolf, M. M., Verstraete, F., Hastings, M. B. & Cirac, J. I. Area laws in quantum systems: Mutual information and correlations. *Phys. Rev. Lett.* **100,** 070502 (2008).
10. Molnar, A., Schuch, N., Verstraete, F. & Cirac, J. I. Approximating Gibbs states of local Hamiltonians efficiently with projected entangled pair states. *Phys. Rev. B* **91,** 045138 (2015).
11. Ge, Y. & Eisert, J. Area laws and approximations of quantum many-body states. Preprint at http://arXiv.org/abs/1411.2995 (2014).
12. White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **69,** 2863–2866 (1992).
13. Schollwöck, U. The density-matrix renormalization group. *Rev. Mod. Phys.* **77,** 259–315 (2005).
14. Hallberg, K. A. New trends in density matrix renormalization. *Adv. Phys.* **55,** 477–526 (2006).
15. White, S. Spin gaps in a frustrated Heisenberg model for CaV_4O_9. *Phys. Rev. Lett.* **77,** 3633–3636 (1996).
16. Hieida, Y. Application of the density matrix renormalization group method to a non-equilibrium problem. *J. Phys. Soc. Jpn* **67,** 369–372 (1998).
17. White, S. & Scalapino, D. Density matrix renormalization group study of the striped phase in the 2d t–J model. *Phys. Rev. Lett.* **80,** 1272–1275 (1998).
18. White, S. & Chernyshev, A. Neél order in square and triangular lattice Heisenberg models. *Phys. Rev. Lett.* **99,** 127004 (2007).
19. Eisert, J. Computational difficulty of global variations in the density matrix renormalization group. *Phys. Rev. Lett.* **97,** 260501 (2006).
20. White, S. R. Density matrix renormalization group algorithms with a single center site. *Phys. Rev. B* **72,** 180403 (2005).
21. Arad, I., Landau, Z. & Vazirani, U. Improved one-dimensional area law for frustration-free systems. *Phys. Rev. B* **85,** 195145 (2012).
22. Kitaev, A. Y., Shen, A. H. & Vyalyi, M. N. *Classical and Quantum Computation* Vol. 47 (American Mathematical Society, 2002).
23. Oliveira, R. & Terhal, B. M. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Inf. Comput.* **8,** 900–924 (2008).
24. Aharonov, D., Gottesman, D., Irani, S. & Kempe, J. The power of quantum systems on a line. *Commun. Math. Phys.* **287,** 41–65 (2009).
25. Gottesman, D. & Irani, S. The quantum and classical complexity of translationally invariant tiling and hamiltonian problems. *Theory Comput.* **9,** 31–116 (2013).
26. Schuch, N., Cirac, I. & Verstraete, F. Computational difficulty of finding matrix product ground states. *Phys. Rev. Lett.* **100,** 250501 (2008).
27. Arad, I., Kitaev, A., Landau, Z. & Vazirani, U. *Proc. 4th Innov. Theor. Comput. Sci.* (*ITCS*) (ACM, 2013).
28. Aharonov, D., Arad, I. & Irani, S. Efficient algorithm for approximating one-dimensional ground states. *Phys. Rev. A* **82,** 012315 (2010).
29. Schuch, N. & Cirac, J. I. Matrix product state and mean-field solutions for one-dimensional systems can be found efficiently. *Phys. Rev. A* **82,** 012314 (2010).
30. Vidal, G. Classical simulation of infinite-size quantum lattice systems in one spatial dimension. *Phys. Rev. Lett.* **98,** 070201 (2007).
31. Corboz, P., Rice, T. M. & Troyer, M. Competing states in the t–j model: Uniform d-wave state versus stripe state. *Phys. Rev. Lett.* **113,** 046402 (2014).
32. Jordan, J., Orús, R., Vidal, G., Verstraete, F. & Cirac, J. I. Classical simulation of infinite-size quantum lattice systems in two spatial dimensions. *Phys. Rev. Lett.* **101,** 250602 (2008).
33. Corboz, P., Orús, R., Bauer, B. & Vidal, G. Simulation of strongly correlated fermions in two spatial dimensions with fermionic projected entangled-pair states. *Phys. Rev. B* **81,** 165104 (2010).
34. Kraus, C. V., Schuch, N., Verstraete, F. & Cirac, J. I. Fermionic projected entangled pair states. *Phys. Rev. A* **81,** 052338 (2010).
35. Barthel, T., Pineda, C. & Eisert, J. Contraction of fermionic operator circuits and the simulation of strongly correlated fermions. *Phys. Rev. A* **80,** 042333 (2009).
36. Pineda, C., Barthel, T. & Eisert, J. Unitary circuits for strongly correlated fermions. *Phys. Rev. A* **81,** 050303 (2010).
37. Shi, Q-Q., Li, S-H., Zhao, J-H. & Zhou, H-Q. Graded projected entangled-pair state representations and an algorithm for translationally invariant strongly correlated electronic systems on infinite-size lattices in two spatial dimensions. Preprint at http://arXiv.org/abs/0907.5520 (2009).
38. Corboz, P. & Vidal, G. Fermionic multiscale entanglement renormalization ansatz. *Phys. Rev. B* **80,** 165129 (2009).
39. Pižorn, I. & Verstraete, F. Fermionic implementation of projected entangled pair states algorithm. *Phys. Rev. B* **81,** 245110 (2010).
40. Verstraete, F., Wolf, M. M., Perez-Garcia, D. & Cirac, J. I. Criticality, the area law, and the computational power of projected entangled pair states. *Phys. Rev. Lett.* **96,** 220601 (2006).
41. Schuch, N., Wolf, M. M., Verstraete, F. & Cirac, J. I. Computational complexity of projected entangled pair states. *Phys. Rev. Lett.* **98,** 140506 (2007).
42. Schuch, N., Pérez-García, D. & Cirac, I. Classifying quantum phases using matrix product states and projected entangled pair states. *Phys. Rev. B* **84,** 165139 (2011).
43. Perez-Garcia, D., Verstraete, F., Wolf, M. M. & Cirac, J. I. Peps as unique ground states of local hamiltonians. *Quantum Inf. Comput.* **8,** 650–663 (2008).
44. Cirac, J. I., Poilblanc, D., Schuch, N. & Verstraete, F. Entanglement spectrum and boundary theories with projected entangled-pair states. *Phys. Rev. B* **83,** 245134 (2011).

Acknowledgements

The first two authors acknowledge support by ARO Grant W911NF-12-1-0541, NSF Grant CCF-0905626 and Templeton Foundation Grant 21674. The third author was supported by the National Science Foundation under Grant No. 0844626 and by the Ministry of Education, Singapore under the Tier 3 grant MOE2012-T3-1-009.

Author contributions

All authors contributed equally to this work.

Additional information

Supplementary information is available in the online version of the paper. Reprints and permissions information is available online at www.nature.com/reprints. Correspondence and requests for materials should be addressed to T.V.

Competing financial interests

The authors declare no competing financial interests.

Methods

Throughout, we consider a Hamiltonian $H=\sum_{i=1}^{n-1} H_i$ acting on n d-dimensional qudits, indexed $1,\ldots,n$ from left to right. Here H_i acts on qudits $\{i,i+1\}$ and satisfies $0\le H_i\le \mathrm{Id}$. We assume that there is a constant gap $\varepsilon:=\varepsilon_1-\varepsilon_0$ between the energy ε_0 of the ground state $|\Gamma\rangle$ and the energy ε_1 of the first excited state.

The canonical basis of $\mathbb{C}^d$ is denoted by $\{|1\rangle,\ldots,|d\rangle\}$. We write $\mathcal{H}$ for the Hilbert space $(\mathbb{C}^d)^{\otimes n}$ corresponding to the n qudits and $\mathcal{H}_{[i,j]}$ for the Hilbert space of the subset of qudits with indices in $[i,j]$; we also write $\mathcal{H}_i$ for $\mathcal{H}_{[i,i]}$. For any density matrix ρ acting on $\mathcal{H}$, $\mathrm{tr}_{[i,j]}\rho$ will denote the tracing out of the qudits with indices in $[i,j]$.

Vectors in $(\mathbb{C}^d)^{\otimes n}$ will be represented as matrix-product states (MPSs), which can be specified as a sequence of tensors $A_1,\ldots,A_n$, where $A_1\in\mathbb{C}^d\times\mathbb{C}^{B_1}$, $A_i\in\mathbb{C}^{B_{i-1}}\times\mathbb{C}^d\times\mathbb{C}^{B_i}$ for $1\le i<n$, and $A_n\in\mathbb{C}^{B_{n-1}}\otimes\mathbb{C}^d$. We will refer to B_i as the bond dimension across cut $(i,i+1)$.

The constant $c_\varepsilon:=(\varepsilon/169)^2$ will play a particular role in our analysis. We note that it satisfies the following inequalities:

$$c_\varepsilon\left(1+\frac{1}{\varepsilon}\right)\le\frac{1}{2},\quad 14\sqrt{c_\varepsilon}/\varepsilon<\frac{1}{12},\quad \frac{84c_\varepsilon}{\varepsilon}<\frac{1}{2} \tag{1}$$

We show the following.

Theorem. Let H be an n-qudit 1D local Hamiltonian with minimal energy ε_0, ground state $|\Gamma\rangle$, gap ε, and $\delta>0$. There is an algorithm that runs in time $n^{c(d,\epsilon_0,\epsilon)}\mathrm{poly}(n/\delta)$, where $c(d,\epsilon_0,\epsilon)=2^{O(\log^3 d/\epsilon)}+O(\varepsilon_0/\varepsilon)$, and with probability at least $1-1/\mathrm{poly}(n)$ returns an MPS representing a state $|\Psi\rangle$ such that $|\langle\Psi|\Gamma\rangle|\ge 1-\delta$.

Algorithm description and analysis. We recall the definition of a viable set.

Definition. Given $\delta>0$ and an integer i, $1\le i\le n$, a set $S\subseteq(\mathbb{C}^d)^{\otimes i}$ is said to be (i,δ)-viable for $|\Gamma\rangle$ if there exists a state $|\phi\rangle\in(\mathbb{C}^d)^{\otimes n}$ such that $|\langle\phi|\Gamma\rangle|\ge 1-\delta$ and such that the reduced density of $|\phi\rangle$ on the first i qudits is supported on $\mathrm{Span}(S)$; we shall call such a state $|\phi\rangle$ a witness for S and δ the error of S.

We will further say that the set S is (i,s,b,δ)-viable for $|\Gamma\rangle$ if $|S|\le s$ and each $v\in S$ can be described by an MPS with maximum bond dimension at most b.

In the ith iteration our algorithm constructs a $(i,s,b,c_\varepsilon/n)$-viable set, where c_ε is a constant depending only on ε that satisfies (1). The four steps in each iteration are designed to update the four parameters of the viable set. As we will show, these updates always satisfy the condition that the parameters s and b are each bounded by some fixed polynomial in n of degree independent of the iteration i.

The initialization step $i=0$ is trivial, as the set $\{1\}$ is a $(0,\delta)$-viable set for any $\delta\ge 0$, $b\ge 0$ and $s\ge 1$. Let $1\le i\le n$ be an integer, and S_{i-1} the $(i-1,s,b,c_\varepsilon/n)$-viable set obtained at the end of the $(i-1)$st iteration of the algorithm, where s and b are both polynomial in n. In the following we describe how each of the four steps of the algorithm can be performed efficiently, and track the changes in the parameters of the viable set.

Extension. The first step in the ith iteration involves extending the set S_{i-1} to an $(i,ds,b,c_\varepsilon/n)$-viable set $S_i^{(1)}$ as follows.

Algorithm step 1: extension

Let S_{i-1} be a $(i-1,s,b,c_\varepsilon/n)$-viable set.

1. Return $S_i^{(1)}:=\left\{|s\rangle|j\rangle:|s\rangle\in S_{i-1},\ 1\le j\le d\right\}$.

The computation of $S_i^{(1)}$ from S_{i-1} can clearly be done efficiently: MPS representations for vectors in the latter are constructed as the concatenation of an MPS for a vector in the former with an independent tensor corresponding to the additional basis state $|j\rangle$. We refer to the Supplementary Methods for a formal proof.

Lemma 1. $S_i^{(1)}$ is an $(i,ds,b,c_\varepsilon/n)$-viable set.

Size and bond trimming. Given $1\le i\le n$, we write $\mathcal{H}_L$ for the space $\mathcal{H}_{[1,i]}$ and $\mathcal{H}_R$ for the space $\mathcal{H}_{[i+1,n]}$. We also let $H_L=H_1+\cdots+H_{i-1}$ and $H_R=H_{i+1}+\cdots+H_{n-1}$, so that the total Hamiltonian $H=H_L+H_i+H_R$, where H_i is the only term acting across the $(i,i+1)$ cut.

Definition. Given a state of Schmidt rank B and Schmidt decomposition across the $(i,i+1)$ cut given by $|v\rangle=\sum_{j=1}^{B}\lambda_j|a_j\rangle|b_j\rangle$, let $U_v:\mathbb{C}^B\to\mathcal{H}_R$ be the partial isometry specified by $U_v|j\rangle=|b_j\rangle$. By abuse of notation we also write U_v^* for $I\otimes U_v^*$ when acting on $\mathcal{H}_{[k,n]}$ for $k\le i+1$.

- Define the left state of $|v\rangle$ to be $\mathrm{ls}(v):=U_v^*|v\rangle=\sum_j\lambda_j|a_j\rangle|j\rangle\in\mathcal{H}_L\otimes\mathbb{C}^B$
- Define the boundary contraction of v as

$$\mathrm{cont}(v):=\mathrm{tr}_{[1,\ldots,i-1]}(|\mathrm{ls}(v)\rangle\langle\mathrm{ls}(v)|)=U_v\,\mathrm{tr}_{[1,\ldots,i-1]}(|v\rangle\langle v|)\,U_v^*$$

Then $\mathrm{cont}(v)$ is a density matrix supported on $\mathcal{H}_i\otimes\mathbb{C}^B$.

We will make use of a η-net over the unit ball of boundary contractions for the trace norm, for $\eta=c_\varepsilon/(2n)$. Such a net can be efficiently constructed by discretizing a region of $\mathbb{C}^B\otimes\mathbb{C}^d$ containing its unit ball, leading to a net $\mathcal{N}=\mathcal{N}_\eta$ of size $|\mathcal{N}|=(Bd/\eta)^{O((Bd)^2)}$. For each $X\in\mathcal{N}$, by solving a suitable convex program we find a state on $\mathcal{H}_L\otimes\mathbb{C}^{B_{c_\varepsilon}}$ of minimum energy among those states whose boundary contraction (reduced density matrix on $\mathcal{H}_i\otimes\mathbb{C}^{B_{c_\varepsilon}}$) is close to X. The new viable set is then the union over all elements of $\mathcal{N}$ of the left Schmidt vectors, on $\mathcal{H}_L$, of these states. Unfortunately, the dimension of this convex program scales with the dimension of $\mathcal{H}_L$, which is exponential in n, and to solve it efficiently we must restrict the optimization to states supported on a subspace S of polynomial dimension.

Algorithm steps 2 and 3: size trimming and bond trimming

Let $S_i^{(1)}$ be the $(i,ds,b,c_\varepsilon/n)$-viable set constructed as a result of the extension step.

2. For each $X\in\mathcal{N}$, solve the following size trimming convex program, whose variable is a density matrix σ supported on the space $\mathrm{Span}\{S_i^{(1)}\}\otimes\mathbb{C}^{B_{c_\varepsilon}}\subseteq\mathcal{H}_L\otimes\mathbb{C}^{B_{c_\varepsilon}}$:

$$\min\ \sum_{j=1}^{i-1}\mathrm{tr}(H_j\sigma) \tag{2}$$

$$\text{such that}\quad \|\mathrm{tr}_{[1,\ldots,i-1]}(\sigma)-X\|_1\le\frac{c_\varepsilon}{2n},$$

$$\mathrm{tr}(\sigma)=1,\ \sigma\ge 0$$

Let $|u\rangle=\sum_j|u_j\rangle|j\rangle$ be the leading eigenvector of the solution σ to this program, and let $S_i^{(2)}$ be the set containing the union of all $\{|u_j\rangle\}$, obtained for each net element X.

3. Trim each of the bonds $1,\ldots,i-1$ of each $|u\rangle\in S_i^{(2)}$ to $p_2(n)$, where $p_2(n)$ is a polynomial defined in Claim 3 below. Include the MPS representation of all resulting vectors in $S_i^{(3)}$.

We note that because the set $S_i^{(1)}$ contains vectors specified using polynomial-size MPSs, for any X a polynomial-size representation for the optimal solution σ to the convex program (2) can be computed efficiently. We refer to the Supplementary Methods for more detail.

The following two lemmas state the properties of the sets $S_i^{(2)}$ and $S_i^{(3)}$. For proofs we refer to the Supplementary Methods.

Lemma 2. If $S_i^{(1)}$ is a $(i,ds,b,(c_\varepsilon/n))$-viable set then $S_i^{(2)}$ is a $(i,p_1(n),dsb,(1/12))$-viable set with $p_1(n)=B_{c_\varepsilon}(4\lceil B_{c_\varepsilon}dn/c_\varepsilon\rceil+1)^{2B_{c_\varepsilon}d}$.

The key property used to bound the error incurred in the trimming step is that the state being trimmed is close to a state with low Schmidt rank, as follows from the strong form of the area law. Let $r(n)$ be a polynomial such that there exists a vector $|v\rangle$ with Schmidt rank $r(n)$ such that $|\langle v|\Gamma\rangle|>1-1/48$ (as is shown to exist in Lemma 5 from the Supplementary Methods).

Lemma 3. The set $S_i^{(3)}$ produced at the end of step 3 of the algorithm is a $(i,p_1(n),p_2(n),1/2)$-viable set, where $p_2(n):=48nr(n)$.

Error Reduction. The final step of the algorithm consists of reducing the error of the viable set, transforming the $(i,p_1(n),p_2(n),(1/2))$-viable set produced as a result of the previous step to a $(i,p(n)p_1(n),p(n)p_2(n),(c_\varepsilon/n))$-viable set, where $p(n)$ is a fixed polynomial.

Algorithm step 4: error reduction

Let $S_i^{(3)}$ be the set constructed as a result of the size and bond trimming steps described in the previous section.

4. Randomly select a sampling AGSP K (defined below) with m and l as in Corollary 1. Decompose K as $K=\sum_j A_j\otimes B_j$. Return $S_i^{(4)}:=\{A_j|s\rangle:|s\rangle\in S_i^{(3)}\}$.

That this step can be carried out in polynomial time follows from the properties of K as detailed below.

Lemma 4. For any choice of polynomial $q(n)$, there exists a polynomial $p(n)$ such that Step 4. described above maps any given $(i,s,b,(1/2))$-viable set S into an $(i,p(n)s,p(n)b,c_\varepsilon/q(n))$-viable set S' with success probability $1-1/n^3$.

The sampling AGSP. Our starting point is an operator A that approximates the projection onto the ground state Γ, defined as

$$A:=\left(\frac{1}{1-\frac{\varepsilon_0}{n}}\left(1-\frac{H}{n}\right)\right)^m$$

The operator K is then formed from a polynomial sample of the exponentially many terms obtained by expanding the mth power in the definition of A. Using a matrix-valued Chernoff bound, this polynomial sample can be shown to provide a close approximation to A with high probability (see Proposition 1). We now proceed with the details.

First note that A is positive semidefinite, has operator norm 1 and satisfies $A|\Gamma\rangle=|\Gamma\rangle$. Furthermore, for any polynomial $q(n)$, fixing $m=\Theta((1/\varepsilon)n\log(q(n)/c_\varepsilon))$ we have that for any $|\Gamma^\perp\rangle$ orthogonal to the ground state $|\Gamma\rangle$ and such that $\||\Gamma^\perp\rangle\|=1$,

$$|\langle\Gamma^\perp|A|\Gamma^\perp\rangle|\le\left(\frac{1-\varepsilon_1/n}{1-\varepsilon_0/n}\right)^m=O\left(\left(1-\frac{\varepsilon}{n}\right)^{\Omega(\frac{1}{\varepsilon}n\log\frac{q(n)}{c_\varepsilon})}\right)\le\frac{c_\varepsilon}{2q(n)}$$

Write $P_i:=(1-H_i)$, $C:=1/(1-\varepsilon_0/n)$. For any integer m, expand

$$A=C^m\frac{1}{n^m}\sum_{I=(i_1,\ldots,i_m)\in\{1,\ldots,n\}^m}P_I,\quad\text{with } P_I:=\prod_{j=1}^{m}P_{i_j} \tag{3}$$

Definition. Define a sampling AGSP operator $K := C^m(1/\ell)\sum_{j=1}^{\ell} P_{l_j}$ to be the average of ℓ terms P_{l_j} chosen uniformly at random from all terms in the expansion (3) of A.

We note that we may not be given ε_0 and therefore cannot specify the constant C explicitly. However, we observe that any multiple of K will suffice for use within the algorithm; only the resulting vectors need to be normalized. The proof of the following is based on a variant of the Matrix-valued Chernoff bound.

Proposition 1. For any polynomial $q(n)$, there exists $m = O((1/\varepsilon)n\log(q(n)/c_\varepsilon))$ and $\ell = n^{O(\varepsilon_0/\varepsilon)}$ (where the implied constants may depend on the degree of $q(n)$) such that with probability at least $1 - 1/n^3$ the sampling AGSP operator K defined above has the following properties: 1. $\|K - A\| \leq (1/q(n))$, 2. Every projection P_j appears no more than $\kappa\log(n)$ times in any term P_l of K for some $\kappa = O(\varepsilon_0/\varepsilon)$.

Proof of main theorem

With the four steps (extension, size trimming, bond trimming, and error reduction) of the algorithm established, the proof of our main theorem follows without difficulty.

Proof of main theorem. Lemmas 1–4 together show that the succession of the four steps of the algorithm detailed in the previous sections transforms any viable set S_{i-1} with parameters $(i-1, p(n)p_1(n), p(n)p_2(n), c_\varepsilon/n)$ into a viable set S_i with parameters $(i, p(n)p_1(n), p(n)p_2(n), c_\varepsilon/n)$, where p, p_1 and p_2 are all fixed polynomials independent of i. Moreover, this transformation can be executed in probabilistic polynomial time, with a success probability at least $1 - (1/n^3)$.

Starting from the set $\{1\}$, which is trivially a $(0, c_\varepsilon/n)$-viable set, and proceeding inductively we efficiently construct an $(n, p(n)p_1(n), p(n)p_2(n), c_\varepsilon/n)$-viable set, with success probability at least $1 - (1/n^2)$. From this viable set we show how to obtain an inverse-polynomial approximation to the ground state.

For this we first observe that the error reduction step in the final iteration can be modified to produce a $(n, p'(n)p_1(n), p'(n)p_2(n), c_\varepsilon/(np(n)))$-viable set S, for any fixed polynomial $p(n)$ of our choice; for this it suffices to set $q(n) = np(n)$ instead of $q(n) = n$ in this step. Note that given the index $i = n$, the condition that S is $(n, c_\varepsilon/(np(n)))$-viable simply means that there is a $c_\varepsilon/(np(n))$ approximation to the ground state supported on S. Such an approximation has energy at most $\varepsilon_0 + (1/p(n))$, and can be found by solving the convex program

$$\min \sum_{j=1}^{n-1} \mathrm{tr}(H_j\sigma)$$
$$\mathrm{tr}(\sigma) = 1,\ \sigma \geq 0$$

which is analogous to (2) but for the omission of the constraint on the boundary contraction. By Lemma 10 from the Supplementary Methods, the leading eigenvector $|u\rangle$ of an optimal solution σ satisfies $|\langle u|\Gamma\rangle| \geq 1 - 1/p(n)$, as required. Moreover, σ and $|u\rangle$ can be computed efficiently, as detailed in the Supplementary Methods.