

Implementation of rigorous renormalization group method for ground space and low-energy states of local Hamiltonians

Brenden Roberts,^{*} Thomas Vidick, and Olexei I. Motrunich*Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, California 91125, USA*

(Received 2 April 2017; published 18 December 2017)

The success of polynomial-time tensor network methods for computing ground states of certain quantum local Hamiltonians has recently been given a sound theoretical basis by Arad *et al.* [*Math. Phys.* **356**, 65 (2017)]. The convergence proof, however, relies on “rigorous renormalization group” (RRG) techniques which differ fundamentally from existing algorithms. We introduce a practical adaptation of the RRG procedure which, while no longer theoretically guaranteed to converge, finds matrix product state ansatz approximations to the ground spaces and low-lying excited spectra of local Hamiltonians in realistic situations. In contrast to other schemes, RRG does not utilize variational methods on tensor networks. Rather, it operates on subsets of the system Hilbert space by constructing approximations to the global ground space in a treelike manner. We evaluate the algorithm numerically, finding similar performance to density matrix renormalization group (DMRG) in the case of a gapped nondegenerate Hamiltonian. Even in challenging situations of criticality, large ground-state degeneracy, or long-range entanglement, RRG remains able to identify candidate states having large overlap with ground and low-energy eigenstates, outperforming DMRG in some cases.

DOI: [10.1103/PhysRevB.96.214203](https://doi.org/10.1103/PhysRevB.96.214203)

I. INTRODUCTION

Many important techniques for solving lattice models in condensed matter physics take the form of tensor network algorithms. The seminal such method is White’s density matrix renormalization group (DMRG) [1], an optimization algorithm on the matrix product state (MPS) ansatz of quantum wave functions [2,3] developed as a controllable method improving on Wilson’s numerical renormalization group for impurity systems [4]. DMRG remains the most versatile procedure in its class. It has been heavily used to numerically solve low-dimensional quantum models; an early instance was the Haldane phase in the Heisenberg chain [5]. More recently, a related approach was used in the classification of all gapped phases in one dimension [6]. Other related techniques include the tensor renormalization group and tensor network renormalization, which utilize a two-dimensional coarse-graining process to solve quantum systems in one dimension by the quantum-to-classical correspondence [7–9]. Other variational algorithms operate on the multiscale entanglement renormalization ansatz (MERA), which efficiently represents states exhibiting logarithmic violation of the area law by encoding correlations at all scales in an optimized quantum circuit of logarithmic depth [10–14].

In this paper, we present an alternative approach to the solution of local Hamiltonians in one dimension (1D). The rigorous renormalization group (RRG) is a recent algorithm developed as part of a proof of the tractability of computing ground states of gapped local Hamiltonians in 1D [15,16]. The proof employs techniques first introduced to establish an improved one-dimensional area law [17]. Broadly, the strategy is as follows: partition the system into small initial blocks and, focusing on the Hilbert space of the blocks individually, identify sets of states that are “extendable” to the rest of the system to create a good approximation to the system-wide ground

space. This property is termed *viability*, and formally defined in Eq. (1). The identification of viable sets is accomplished with an *approximate ground-state projector* (AGSP), an operator approximately filtering out highly excited states on the entire system, whose support is restricted to perform this filtering within each block individually. In this way RRG deviates from a traditional real-space blocking scheme, in which each block does not have access to global information. The next step is to merge the identified viable sets on adjacent blocks, obtaining states supported on blocks of larger size. However, this step and the local application of the AGSP result in an untenable blowup of the number of states, so a reduction step is performed, returning the number of states per block (now comprising two blocks of the smaller size) to a constant value. This procedure is iterated, merging blocks in a treelike manner, and at the full system scale, the identified states are shown to closely approximate the low-energy space [15].

In the present work we adapt these techniques to specify a concrete RRG procedure allowing for the explicit computation of ground and low-energy states of local Hamiltonians. This requires making allowance for computational limitations, and generally our modifications operate outside of the regime of rigorous guarantee. Still, our algorithm presents a conceptually distinct approach to this task. We emphasize that the use of the word “rigorous” is in reference to the title of Arad *et al.* [15], rather than in order to establish a contrast with other tensor network algorithms.

The main conceptual departure of this algorithm from existing tensor network methods is that RRG operates on viable sets of states supported on blocks, rather than on variational states in the full Hilbert space. Two important features arise from this distinction. First, no local energy minimization on a particular ansatz state is performed. Even though in the RRG procedure described here the basic operations are performed on MPS comprising an approximate basis of the viable sets, the MPS objects themselves are incidental, and the concerns arising from the MPS ansatz (e.g., gauge choice, truncation) are external to the fundamental algorithm.

^{*}broberts@caltech.edu

- The input is a local Hamiltonian H acting on N qubits, specified by an MPO. Let n , s and D be input parameters.
1. Initialize:
 - (a) Construct approximate ground state projector (AGSP) K from Hamiltonian H
 - (b) Partition system into contiguous blocks of length n , denoted J_0^λ , $\lambda = 0, 1, \dots, N/n - 1$. Obtain s -dimensional low-energy eigenspace V_0^λ of block Hamiltonian H_0^λ for each λ .
 2. For $m = 0, 1, \dots, \log_2(N/n) - 1$, denoting an ‘‘RG step’’ or scale factor:
 - (a) Expand: for $\lambda = 0, 1, \dots, N/(2^m n) - 1$:
 - i. Extract D^2 operators $\{A_{m,r}^\lambda\}_{r=1, \dots, D^2}$ from the AGSP K , acting on subsystem J_m^λ . Operate on the viable set, taking $V_m^\lambda \rightarrow W_m^\lambda \equiv \{A_{m,r}^\lambda V_m^\lambda\}_r$, where $\dim(W_m^\lambda) \leq sD^2$.
 - ii. Compute the restriction of the block Hamiltonian H_m^λ to $W_m^\lambda \subset \mathfrak{H}_m^\lambda$.
 - (b) Reduce: For $\lambda = 0, 2, \dots, N/(2^m n) - 2$:
 - i. (Merge) Obtain the tensor product space $W_m^\lambda \otimes W_m^{\lambda+1} \subset \mathfrak{H}_{m+1}^{\lambda/2}$, supported on qubits in $J_{m+1}^{\lambda/2} = J_m^\lambda \cup J_m^{\lambda+1}$. Compute the restriction of $H_{m+1}^{\lambda/2}$ to the tensor product set.
 - ii. Obtain s -dimensional low-energy eigenspace of the restriction of $H_{m+1}^{\lambda/2}$ to the tensor product space. Use the eigenstates as a basis for viable set $V_{m+1}^{\lambda/2}$ in iteration $m + 1$.
 3. At $m = m^* = \log(N/n)$, the viable set $V_{m^*}^0$ is a candidate for the low-energy space T supported on the full system.

FIG. 1. Outline of the implemented RRG algorithm.

Second, the physical degrees of freedom are not coarse-grained. The objective of a coarse-graining strategy is to limit the dimensionality of the Hilbert space at increasing scale by the introduction of renormalized degrees of freedom, determined by some local rule, specifying a smaller effective Hilbert space. Instead, RRG achieves this goal by maintaining viable sets of constant dimension at all levels of the algorithm hierarchy. These processes cannot be considered equivalent, as the RRG step of applying the AGSP operator changes the relationship between scales in a complicated way, and does not match the intuition of an ‘‘RG flow’’ in a small number of parameters. However, this method still allows for fully controllable systematic improvements in accuracy.

The structure of this paper is as follows. We first give a detailed, self-contained description of our algorithm in Sec. II. (We refer the reader familiar with the theoretical RRG paper to the Appendix for a precise discussion of the differences between the proof and the present work.) We provide an extended discussion of numerical results in Sec. III. Given its origins as a highly technical theoretical algorithm developed in order to obtain provable guarantees, the RRG method performs surprisingly well, often matching the results of standard DMRG implementations and outperforming them in certain difficult cases exhibiting degenerate ground spaces or highly entangled ground states. Finally, we conclude and give directions for further work in Sec. IV.

II. OPERATION OF ALGORITHM

A. Overview and notation

The steps of RRG as implemented are listed in Fig. 1 for reference and are discussed in detail in subsequent sections. A visual schematic is shown in Fig. 2. Our notation is as follows. Let $H = \sum_{i=0}^{N-2} h_i$ be a 2-local Hamiltonian on a chain of N

qubits, with term h_i acting on sites i and $i + 1$. (The generalization to k -local Hamiltonians and qudits is straightforward.) Denote the Hilbert space of the system by \mathfrak{H} , and refer to the low-energy eigenspace of H as T . Let n be a parameter specifying the size of initial regions of the system, and assume N/n is a power of 2. For each $m = 0, 1, \dots, \log_2(N/n)$, partition the N -site system into contiguous blocks of equal length $2^m n$. Call these $J_m^\lambda = \{\lambda 2^m n, \dots, (\lambda + 1)2^m n - 1\}$, for $\lambda = 0, 1, \dots, N/(2^m n) - 1$. The Hilbert space associated with J_m^λ is denoted \mathfrak{H}_m^λ , and $\mathfrak{H} = \bigotimes_\lambda \mathfrak{H}_m^\lambda$. Let H_m^λ be the block Hamiltonian on J_m^λ , comprising all terms acting only on sites in J_m^λ and excluding boundary terms. Explicitly, $H_m^\lambda = \sum_{i \in J_m^{\lambda*}} h_i$, where $J_m^{\lambda*} = \{\lambda 2^m n, \dots, (\lambda + 1)2^m n - 2\}$.

B. Initialization

The first step is to construct an AGSP K , whose action on states in \mathfrak{H} increases overlap with T , the low-energy subspace of H . Many constructions of AGSP are possible. In the interest of efficiency, we use an AGSP obtained as an approximation to a thermal operator at temperature t/k , $K \approx e^{-kH/t}$, $t, k > 0$. Let Q_t denote a matrix product operator (MPO) approximating the thermal operator $e^{-H/t}$ at temperature t ; procedures such as a Trotter decomposition [18] or cluster expansion can be used to efficiently compute this MPO. The AGSP is then obtained as a power of Q_t , contracting the product on the physical indices k times.

Because the AGSP must later be divided into operators acting on individual blocks, to compute Q_t requires contraction of the tensor network having terms of the form $e^{-h_i/t}$. After each contraction a singular value decomposition (SVD) is performed between site indices, and the MPO is truncated by eliminating low-weight Schmidt vectors across each bond. Here truncation is meant in the sense of MPS truncation, representing the MPO as a state in a higher-dimensional local

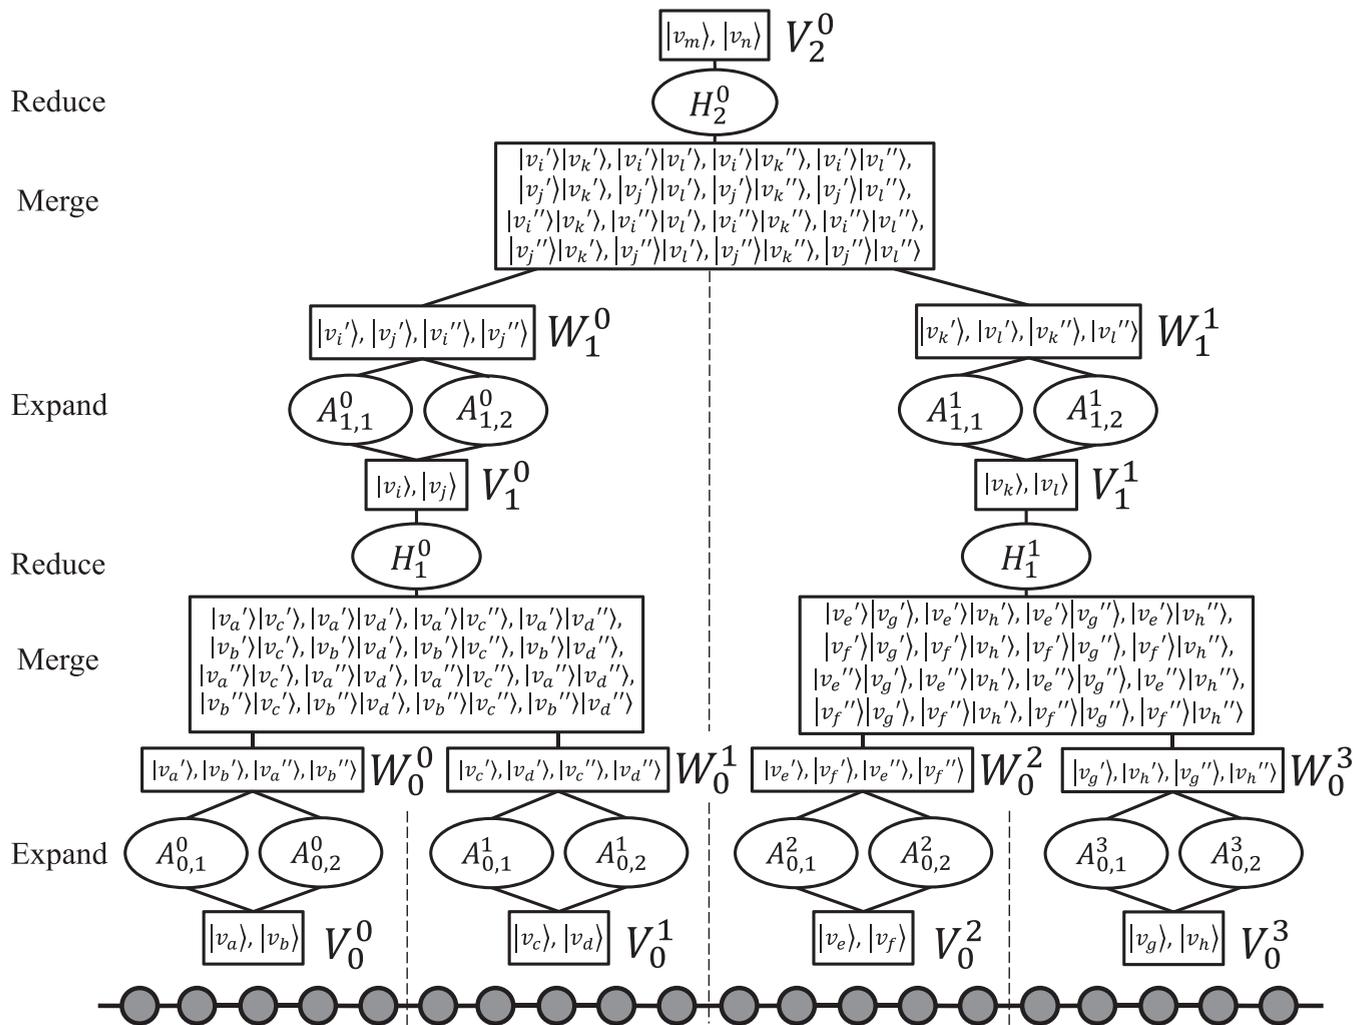


FIG. 2. Schematic illustration of the RRG algorithm over several length scales $m = 0, 1, 2$. As shown, $(s, D) = (2, 2)$ and block size $n = 5$. Gray dots represent local Hilbert spaces, and the V_m^0, V_m^1 , etc. are the viable sets over blocks of sites. The labeled vectors $|v_a\rangle$, etc., are basis states for the viable sets and have no relationship between blocks at a given scale. The action of the projector operators $A_{m,r}^\lambda$ on the states is represented by primes and double primes (e.g., $A_{0,1}^\lambda |v_a\rangle = |v_a'\rangle$) and $A_{0,2}^\lambda |v_a\rangle = |v_a''\rangle$). These generate the expanded viable sets W_m^0 , and so on. The merge procedure obtains tensor product states such as $|v_a'\rangle|v_c'\rangle$ supported on two blocks, and the tensor product set is reduced in dimension via diagonalizing block Hamiltonians like H_1^0 , producing a viable set supported on two of the previous blocks.

Hilbert space. This amounts to using the Frobenius norm to order the terms arising from the SVD, and may not be an optimal way to approximate operators; we address this issue in more detail later.

The second step in the initialization is to identify sets of states $V_0^\lambda \subset \mathfrak{H}_0^\lambda$, for $\lambda = 0, 1, \dots, N/n - 1$, of constant dimension s , where s is a parameter of the algorithm which bounds the dimension of the sets manipulated throughout. We use the term “viable sets” for the V_0^λ (generally, for $V_m^\lambda \subset \mathfrak{H}_m^\lambda$) because the intent of the algorithm is that each V_m^λ be extendable to include a good approximation to the global low-energy eigenspace T . That is, each set V_m^λ is chosen such that if $\mathfrak{H} = \mathfrak{H}_m^\lambda \otimes \bar{\mathfrak{H}}_m^\lambda$, then $V_m^\lambda \otimes \bar{\mathfrak{H}}_m^\lambda$ contains a subspace which is a good approximation to T . We identify a set V_m^λ as δ viable if

$$P_T P_{V_m^\lambda \otimes \bar{\mathfrak{H}}_m^\lambda} P_T \geq (1 - \delta) P_T, \quad (1)$$

where P_T is a projector onto a subspace T . More concretely, consider the case of a nondegenerate global ground space $T = \text{Span}\{|\tau\rangle\}$, $|\tau\rangle \in \mathfrak{H}$. The viability of the set V_m^λ is given by

$$\delta = 1 - \max_{|x\rangle \in V_m^\lambda \otimes \bar{\mathfrak{H}}_m^\lambda} |\langle \tau | x \rangle|^2, \quad (2)$$

where $|x\rangle = \sum_j a_j |v_j\rangle |\bar{v}_j\rangle$ for a collection of states $\{|v_j\rangle\} \subset V_m^\lambda$ along with coefficients a_j , and states $\{|\bar{v}_j\rangle\}$ arbitrary in the Hilbert space of the sites in the complement. It will be shown in Sec. II C that one need never explicitly compute the $\{|\bar{v}_j\rangle\}$. For the case that $\dim(T) > 1$, δ is obtained by taking the smallest value of the maximum in (2), over all $|\tau\rangle \in T$. The goal of the algorithm is to construct the viable sets V_m^λ in such a way that they are indeed δ viable for some constant δ less than 1 for all scales m . Note that a small value of δ corresponds to a better approximation, in contrast with measures like overlap. We emphasize that the viability parameter is not explicitly

computed by the algorithm. Instead, it provides a useful metric to evaluate performance, both in terms of the theoretical results and in terms of experimental investigations for cases where we wish to compare with other methods providing estimates for the ground space (such as when exact diagonalization is possible).

If n is chosen to be small enough, generic operators on \mathfrak{H}_0^λ can be exactly diagonalized. In the initialization step, the initial viable set V_0^λ is specified to be the span of the s eigenvectors of H_0^λ of lowest energy, obtained by exact diagonalization.

C. Iteration over scale

The algorithm proceeds through a treelike hierarchy, the levels of which are specified by a scale parameter $m = 0, 1, \dots, \log_2(N/n)$. At scale m , block J_m^λ consists of $2^{2^m n}$ sites and the region index λ runs from 0 to $N/(2^{2^m n}) - 1$. Note that although the scale of the algorithm is increasing, we do not eliminate any of the physical degrees of freedom. At each step we assume that the previous level has produced a viable set V_m^λ with basis $\{|v_q\rangle\}_{q=1,\dots,s}$ represented by MPS, for every λ .

The algorithm performs two steps. The first step is the *expansion* of the viable set, which has the effect of improving the viability parameter δ as defined in Eq. (1). This is accomplished using the AGSP constructed in the initialization step as follows. Let $J_{m,L}^\lambda$ denote the qubits to the left of J_m^λ , and $J_{m,R}^\lambda$ those to the right. (Generally J_m^λ has two boundaries with its complement $J_{m,L}^\lambda \cup J_{m,R}^\lambda$. The system-edge cases follow immediately.) Consider the MPO representation of the AGSP K , whose elementary tensors are collections of operators on the local Hilbert space, as an MPS. The Schmidt decomposition of K across the left boundary, separating $J_{m,L}^\lambda$ from $J_m^\lambda \cup J_{m,R}^\lambda$, produces a virtual index of dimension ζ :

$$K = \sum_{\alpha < \zeta} \sigma_\alpha L_\alpha M_\alpha. \quad (3)$$

The L_α are the left Schmidt vectors and the M_α the right—which are operators on $J_m^\lambda \cup J_{m,R}^\lambda$ —each with a corresponding Schmidt coefficient σ_α . The Schmidt decomposition may then be obtained for each of the M_i across the boundary between J_m^λ and $J_{m,R}^\lambda$, producing a virtual index of dimension ξ . That is,

$$M_\alpha = \sum_{\beta < \xi} v_{\alpha\beta} A_{\alpha\beta} R_{\alpha\beta}. \quad (4)$$

Each $A_{\alpha\beta}$ is an operator on \mathfrak{H}_m^λ , with weight $\gamma_{\alpha\beta} = \sigma_\alpha v_{\alpha\beta}$ in the expansion of K . For clarity we make the algorithm variables explicit: $A_{m,\alpha\beta}^\lambda$. Now let $D > 0$ be another parameter of the algorithm. In order to increase the viability of the set V_m^λ , act on it with the D^2 operators $A_{m,r}^\lambda$, $r = (\alpha, \beta)$, having highest weight $\gamma_r = \gamma_{\alpha\beta}$. That is, take $V_m^\lambda \rightarrow W_m^\lambda = \text{Span}(\{A_{m,r}^\lambda |v_q\rangle\}_{r,q})$, which we refer to as an *expanded viable set* with dimension bounded by sD^2 .

One expects this operation to produce a set W_m^λ of better viability than V_m^λ because the $A_{m,r}^\lambda$ operators together are meant to increase overlap with the global low-energy space T : this is the defining property of the AGSP. More precisely, let $\{|v_j\rangle\}$ be a collection of states in V_m^λ such that there exists $\{|\bar{v}_j\rangle\} \in \bar{\mathfrak{H}}_m^\lambda$ such that for some coefficients a_j , the state $|x\rangle =$

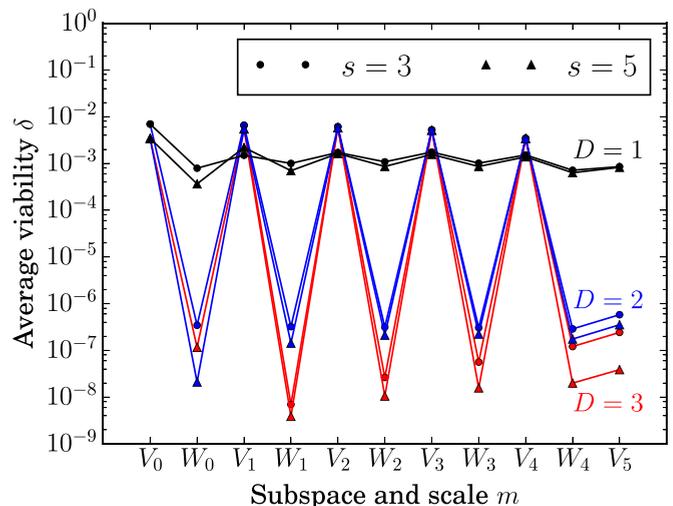


FIG. 3. Viability of sets V_m^λ, W_m^λ averaged over λ , for the non-integrable Ising model with $N = 256$ spins, obtained as the RRG algorithm progresses through the scale hierarchy. Data are shown for parameter values $s = 3, 5$ and $D = 1, 2, 3$.

$\sum_j a_j |v_j\rangle |\bar{v}_j\rangle$ has good overlap with T . By construction, $K|x\rangle$ has better overlap with T than $|x\rangle$. Using the decomposition of Eqs. (3) and (4),

$$K|x\rangle = \sum_{\alpha, \beta, j} \gamma_{m,\alpha\beta}^\lambda a_j A_{m,\alpha\beta}^\lambda |v_j\rangle \otimes \bar{A}_{m,\alpha\beta}^\lambda |\bar{v}_j\rangle, \quad (5)$$

where $\bar{A}_{m,\alpha\beta}^\lambda = L_{m,\alpha}^\lambda R_{m,\alpha\beta}^\lambda$. In this way the viability as defined in Eq. (2) of the set V_m^λ can be improved while leaving both the states and the operators supported on the complement $\bar{\mathfrak{H}}_m^\lambda$ entirely implicit.

If all operators $A_{m,\alpha\beta}^\lambda$ were applied to V_m^λ , the resulting set would contain the collection of states $\{A_{m,\alpha\beta}^\lambda |v_j\rangle\}$, which has improved viability. However, instead of applying all $A_{m,\alpha\beta}^\lambda$, which would lead to an unmanageable blowup in the size of the viable set, we introduce an approximation by selecting the D^2 operators $A_{m,r}^\lambda$ of highest weight γ_r in order to obtain W_m^λ . There is no formal guarantee that this is the best choice, as the Schmidt decomposition is based on the Frobenius rather than the operator norm. In practice we found the choice to be quite reasonable: to observe the increase in viability in a nondegenerate gapped model, compare the V and W points in Fig. 3, and in a critical model in Figs. 5, 6.

The second step performed at each scale m is that of *reduction* of the dimension of the expanded viable sets W_m^λ and W_{m+1}^λ to generate $V_{m+1}^{\lambda/2}$. At the cost of a loss of viability, this step restores s -dimensionality, resulting in a viable set suitable to use at the next level. One first performs a *merge* operation on disjoint pairs of blocks $(\lambda, \lambda + 1)$, with $\lambda = 0, 2, \dots, N/(2^{2^m n}) - 2$. Merging refers to computing the tensor product set $W_m^\lambda \otimes W_{m+1}^{\lambda+1}$ that has support on sites $J_m^\lambda \cup J_{m+1}^{\lambda+1}$. One obtains the viable set $V_{m+1}^{\lambda/2}$, a subspace of $\bar{\mathfrak{H}}_{m+1}^{\lambda/2} = \bar{\mathfrak{H}}_m^\lambda \otimes \bar{\mathfrak{H}}_{m+1}^{\lambda+1}$, from the s -dimensional low-energy eigenspace of the restriction of $H_{m+1}^{\lambda/2}$ to $W_m^\lambda \otimes W_{m+1}^{\lambda+1}$. We note

that this step differs from its counterpart in the theoretical algorithm, which proceeds via random sampling instead of deterministically selecting the lowest-energy eigenvectors of $H_{m+1}^{\lambda/2}$, as we do here. Our choice is based on efficiency considerations described below; see also the Appendix for further discussion. The effect of the operation on the viability of the reduced subspaces can be seen in Figs. 3, 5, and 6.

The single viable set $V_{m^*}^0$ generated at $m^* = \log_2(N/n)$ after the reduction step at scale $m^* - 1$, is a constant-dimensional δ -viable subspace with support on the full system. The algorithm returns the s lowest-energy eigenvectors of the restriction of H to $W_{m^*-1}^0 \otimes W_{m^*-1}^1$, which comprise a basis for this candidate subspace.

D. Scaling and computational considerations

The accuracy with which RRG approximates low-energy eigenstates of H is controlled primarily by two parameters, s and D . To recapitulate, s bounds the dimension of the reduced viable sets at each step, and D controls the level of approximation in the application of the AGSP via the operators $\{A_{m,r}^\lambda\}$, $r = 1, \dots, D^2$. Both parameters are reflected in the bound on the dimension sD^2 of the expanded viable sets W_m^λ .

We review the steps in the algorithm and discuss their complexity scaling based on these parameters. In addition to s and D , important parameters are the system size N and the bond dimensions χ for MPS and η for MPO that are manipulated throughout. For physical Hamiltonians it is reasonable to expect χ and η to be constant in the gapped case, and in gapless systems $\chi, \eta \sim N$. See Schollwöck [3] for a discussion of the scaling of basic MPS operations. Note that the initial block size n only enters this analysis in determining the number of necessary layers $\log(N/n)$.

The initialization requires obtaining viable sets V_0^λ of the Hilbert space \mathfrak{H}_0^λ on the qubits J_0^λ . For small enough choices of n the complexity of this step will be negligible, so we omit it. Similarly, the computation of the full AGSP $K \approx (e^{-H/t})^k$ can be done efficiently via Trotter decomposition, and is not an important bottleneck. In order to extract the operators $\{A_{m,r}^\lambda\}$, $r = 1, \dots, D^2$, the AGSP must be obtained as an MPO in canonical form, analogous to that used for MPS. To do so requires a sequence of $\mathcal{O}(N)$ SVD operations, each with cost $\mathcal{O}(\eta^3)$.

For the steps comprising the iterated procedure we give scaling results applicable at the final computational level $m = m^* - 1$. The first step is to apply K to each V_m^λ by means of the Schmidt decomposition of K across the boundary separating J_m^λ from its complement $\bigcup_{\lambda' \neq \lambda} J_m^{\lambda'}$. This yields a set of operators acting on \mathfrak{H}_m^λ . Applying the D^2 such operators of highest Schmidt weight to a basis of the subspace takes $V_m^\lambda \rightarrow W_m^\lambda$, increasing the dimension to sD^2 . The total cost of contracting these MPSs and MPOs is $\mathcal{O}(sD^2N\chi^2\eta^2)$.

The second step acts on disjoint pairs of neighboring regions, forming the tensor product of expanded viable sets: $W_m^\lambda \otimes W_m^{\lambda+1}$, with dimension $(sD^2)^2$. We compute the matrix elements of the restriction of the block Hamiltonian to the tensor product set. The scaling of this step is $\mathcal{O}([(sD^2)^2]^2N\chi^3)$. For local Hamiltonians the constant can be improved using the

decomposition

$$\begin{aligned} H_{m+1}^{\lambda/2} &= H_m^\lambda + H_m^{\lambda+1} + B_{m+1}^{\lambda/2} \\ &= H_m^\lambda + H_m^{\lambda+1} + \sum_p B_{m,p}^\lambda \otimes B_{m,p}^{\lambda+1}. \end{aligned} \quad (6)$$

The operator $B_{m+1}^{\lambda/2}$ contains $\mathcal{O}(1)$ terms in H acting across the boundary between J_m^λ and $J_m^{\lambda+1}$.

Exact diagonalization of the restricted block Hamiltonian in the subspace has complexity $\mathcal{O}([(sD^2)^2]^3) = \mathcal{O}(s^6D^{12})$. After this, the final step is to explicitly compute the s lowest-energy eigenstates, which has a total cost $\mathcal{O}[s(sD^2)^2N\chi^3]$. These states are used as a basis for the viable set at the next iteration.

From this coarse analysis it is clear that the limiting step with respect to s and D is the diagonalization of the restricted block Hamiltonian. This step is not part of the original formulation, which specifies instead that the reduction of viable set dimension takes place by randomly selecting states from the tensor product set. The choice of our variant is motivated by its effect on the entanglement of the intermediate basis states: low-energy excited states of a block Hamiltonian may display lower entanglement than states chosen randomly. In practice this lowers χ in some systems. It also demonstrates a different possible interpretation of the parameter s , which during the iteration step implicitly defines an energy scale with respect to the restricted Hamiltonian. States having block excitation energy higher than this scale are inaccessible to the algorithm for the purposes of the expansion step.

III. NUMERICAL RESULTS

We now present results from RRG for some example models with the following goals in mind. We first validate the algorithm in a simple gapped nondegenerate system in Sec. III A, demonstrating consistency with DMRG as well as previous numerical and perturbation theory results. In this case the states obtained by RRG are of similar accuracy to those of DMRG, with run times a factor of 5–10 slower depending on s , D , and N . However, we emphasize that it is not the objective of RRG to obtain a numerically precise ground state; rather, it is to accurately identify states having constant overlap with the global low-energy subspace. One expects an optimization algorithm to obtain a more precise state in the absence of local energy minima or very flat energy landscapes, and for simple models we take the DMRG ground state to be exact (in particular, using it to measure viability δ). The RRG candidate states may later be variationally optimized in order to achieve a particular accuracy, but we do not modify the states here.

Our next goal is to demonstrate the practical scaling of the algorithm's performance and computational costs associated with the subspace parameters (s, D). We use the familiar case of the Ising model in the transverse field in Sec. III B, both away from and at criticality. We find that for low values of these parameters, often surprisingly good results can be obtained, with close to unity overlap between DMRG and RRG ground-state candidates. However, neither algorithm scales linearly with system size in the critical regime. Here the slowdown of RRG is no longer a simple numerical factor

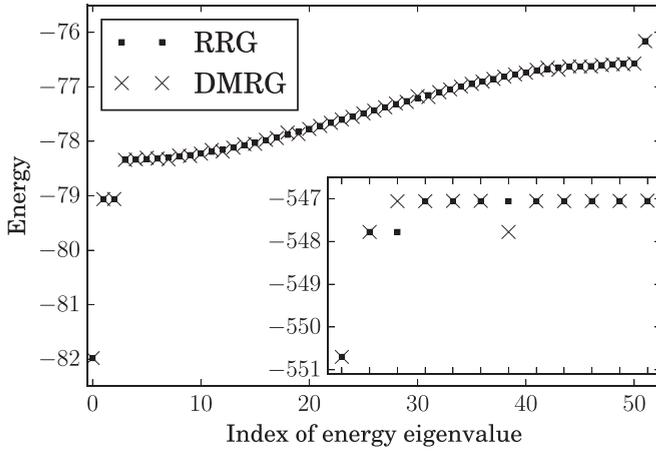


FIG. 4. Energy eigenvalues of the nonintegrable Ising Hamiltonian for $N = 48$ within the subspace obtained by RRG for $(s, D) = (52, 3)$, along with DMRG results for low-energy states. Inset: the same computation for $N = 320$ and $(s, D) = (12, 3)$. DMRG does not consistently identify both edge states in sequence (see text for details).

but becomes a significant cost at larger system sizes (beyond a few hundred sites in our implementation) or for larger values of the algorithm parameters.

Finally, we consider somewhat more challenging models demonstrating areas in which RRG may hold an advantage. In Sec. III C we investigate the Bravyi-Gosset model [19], which has $\mathcal{O}(N)$ ground-state degeneracy, by obtaining a complete basis for the ground space. In Sec. III D we consider the XY model with randomly distributed couplings. The ground state of this model, the random singlet phase, displays long-range entanglement in that it supports algebraic decay of correlations. We compare the correlations present in the candidate states of DMRG and RRG to exact results obtained by the Jordan-Wigner transformation, finding that RRG more accurately reproduces observables measured on the state.

All numerical results were obtained using the tensor network library ITensor [20] for both the DMRG and RRG computations. In all of the following, a Trotter decomposition with 60 steps was used to obtain the tensor network for $Q_t \approx e^{-H/t}$, with $t = 10$, and degree $k = 8$ used to compute the AGSP $K \approx (Q_t)^k$. Thus the effective temperature t/k is of order unity. For reasonable choices of parameters the accuracy of the approximation Q_t is not a limiting factor of the algorithm. Computations were performed on standard hardware on a single node of a computing cluster, with only single threading for the reported run times. A single error parameter τ was used to control MPS truncation in ITensor for both DMRG and RRG (usually $\tau \sim 10^{-9} - 10^{-12}$); in most cases a more lenient value would drastically improve run times with little effect on accuracy. DMRG convergence was handled using a fixed number of sweeps ≥ 20 and relying on the internal diagonalization routine included in ITensor without any modifications specific to the individual systems. Excited states were found iteratively in DMRG by adding projectors into previously found states to the Hamiltonian and using random trial wave functions. Often the average viability

will be used as a metric; this is simply the average over region label λ of the viability δ of each viable set (V_m^λ or W_m^λ) at fixed level m .

A. Nonintegrable Ising model

This model refers to a spin-1/2 Hamiltonian

$$H = -J \sum_{i=0}^{N-2} \sigma_i^z \sigma_{i+1}^z - g \sum_{i=0}^{N-1} \sigma_i^x - h \sum_{i=0}^{N-1} \sigma_i^z. \quad (7)$$

For $h \neq 0$ the model is gapped with a nondegenerate ground state, and admits no good quantum numbers due to the longitudinal component of the field. A recent numerical study [21] for the parameters $(J, g, h) = (1, -1.05, 0.5)$ found the ground-state energy density to be $\varepsilon_0/N \approx -1.722$ and the gap $\gamma = 3.6401$.

We run the RRG algorithm for a fixed system size $N = 256$, initial block size $n = 8$, and track the average viability δ of the viable sets V_m and W_m through the sequence of dimensional expansion and reduction at each scale m (see Fig. 2). Each data point shown in Fig. 3 is the average over λ at a given length scale m . The parameters (s, D) are varied to demonstrate their influence on the results. For gapped systems of this size both DMRG and RRG have run times scaling linearly with system size; however, RRG runs more slowly by a factor of 5–10 compared to DMRG. At $N = 256$, DMRG took 5 min to converge $s = 5$ states (ground and four excited) and RRG ran in 30 min with $(s, D) = (5, 3)$.

The large improvement in viability from V to W is attributable to the AGSP, rather than a simple increase in dimension. Both $\dim(V) = s$ and $\dim(W) = sD^2$ are constant in m and very small compared to the dimensions of the block Hilbert spaces. Choosing n vectors without bias from an M -dimensional space will produce a subspace whose squared overlap with a specific vector is of order n/M . Since M here is exponentially large, a constant increase in n would not much affect measured viability. Thus, the AGSP is an effective projector even at low values of D , which we expect as the model is gapped.

A consequence is that the accuracy of RRG for the largest (s, D) is comparable to that of DMRG, but we do not expect this to be a general feature. Recall from Sec. II B that the criterion the algorithm seeks to maintain is that the measured viability δ of the V_m^λ (and thus the average viability) be bounded for all m by some constant $\delta^* < 1$, rather than approaching unity exponentially in m . The viability of the W_m^λ is not necessarily specified, but should be sufficiently good for the V_{m+1} viable sets at the next level to satisfy the bound. For assessing the performance of RRG, as in Fig. 3, one seeks that δ be maintained away from 1 for the V_m averages.

The final s -dimensional viable sets V_{m^*} (V_5 in Fig. 3) here and in the following examples display much better average viability than that of the previous V_m . This is generally true: at steps $m < m^*$ the viable set is found by diagonalizing a block Hamiltonian H_m^λ , which omits terms present in H . The low-energy eigenspace of this operator need not be close to T , the global low-energy space. At $m = m^*$, however, the low-energy eigenspace of $H_{m^*}^0 = H$ coincides with T ,

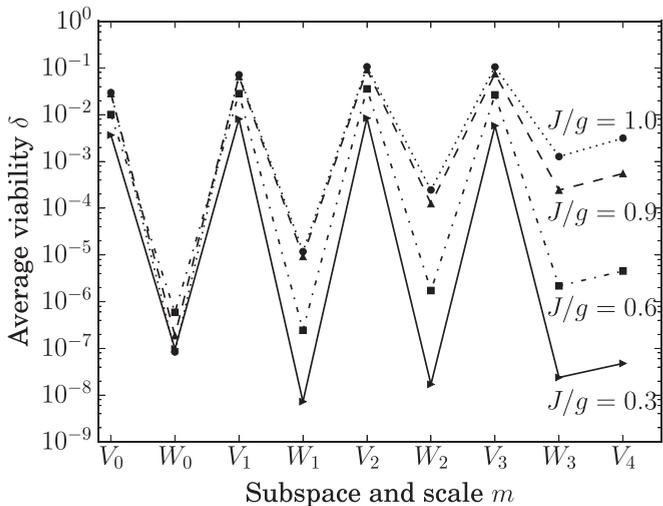


FIG. 5. Viability of sets V_m^λ, W_m^λ , averaged over λ , for the transverse-field Ising model both away from and at criticality. The number of spins is $N = 128$. All data points were generated using parameter values $(s, D) = (5, 4)$.

resulting in minimal loss of viability from the dimensional reduction.

By changing the parameters of RRG, we obtain candidates for low-energy excited states. The ground state of this model is close to a uniform spin-up state, and the excited band contains a spin-flip excitation. Under open boundary conditions two nearly degenerate lower-energy states separate from the first band, corresponding to quasiparticles localized at either edge. We obtain the low-energy spectrum for $N = 48$ with $(s, D) = (52, 3)$, and for $N = 320$ with $(s, D) = (12, 3)$. The results are shown in Fig. 4, compared with DMRG states. For small N both methods find the entire first excited band. In the larger system, the localized edge states are more difficult for DMRG, and it does not consistently find the edge states in sequence. The RRG ground-state energy density at $N = 320$ is $\varepsilon_0/N = -1.721$ and the gap to the excited band is $\gamma = 3.6402$, in agreement with previous results. We find the half-chain entanglement entropy of the ground state and edge states to be $S = 0.01$ bits, and of the states in the band to be $S \approx 1.01$ bits, consistent with qualitative understanding of these states. For DMRG and RRG, ground states have bond dimension 4 and excited states in the band have bond dimension 31. (The methods do not yield identical bond dimension in all cases.)

B. Transverse-field Ising model

Consider the same Hamiltonian in the regime $h = 0$; that is, the Ising model in a transverse field. Figure 5 shows the result as we approach the critical point $J = g$ from the paramagnetic phase for $N = 128$, measuring average viability throughout the algorithm. One observes a strong deterioration of the measured viability as the gap closes. Approaching the critical point, RRG takes increasingly more time than DMRG to run: runtimes for $J/g = 0.6$ for both methods are shown in

TABLE I. Runtimes of DMRG and RRG for the transverse-field Ising model with $J/g = 0.6$, using $(s, D) = (5, 4)$. Some randomness is inherent in the DMRG results due to the use of random trial states. $s = 5$ states are found by DMRG.

N	RRG runtime (s)	DMRG runtime (s)
32	158	94
48	337	132
64	866	208
96	1871	277
128	3912	393

Table I, whereas, for example, at criticality DMRG takes 800 s and RRG takes 17 000 s.

We demonstrate the scaling with parameters s and D at criticality in Fig. 6. The improvement in viability with increasing D is less dramatic than seen in Fig. 3, corresponding to a flatter spectrum of Schmidt values across the cuts between subsystems. Note in this case that at the critical point, as the algorithm progresses, the average viability of the V_m sets visibly approaches unity, in contrast to the gapped case, which appears to maintain viability bounded away from 1.

C. Bravyi-Gosset model

This model was initially introduced as a classification scheme for frustration-free 2-local Hamiltonians [19]. The Hamiltonian is

$$H = \sum_{i=0}^{N-2} |\psi\rangle\langle\psi|_{i,i+1}, \quad (8)$$

where $|\psi\rangle$ is a generic state on two qubits. Up to a global phase, such a state can be specified in the form $|\psi\rangle = R(\theta)_1(p|00\rangle + \sqrt{1-p^2}|11\rangle)$, with $R(\theta)_1$ a rotation performed on the first

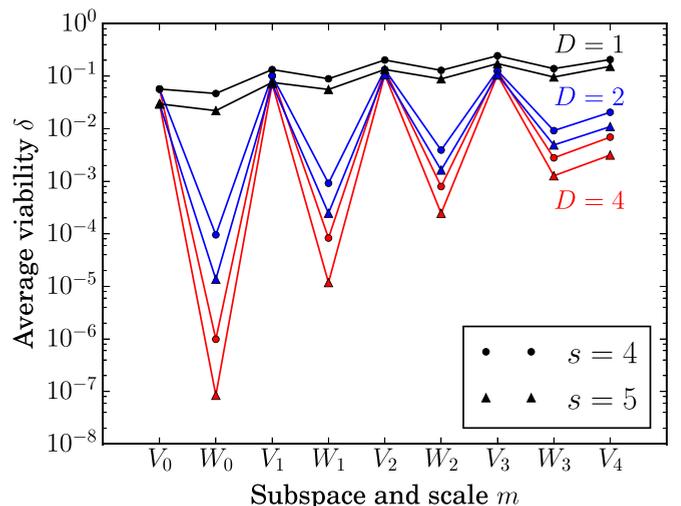


FIG. 6. Viability of sets V_m^λ, W_m^λ , averaged over λ , for the transverse-field Ising model at criticality, obtained as the RRG algorithm progresses through the scale hierarchy. Data are shown for parameter values $s = 4, 5$ and $D = 1, 2, 4$.

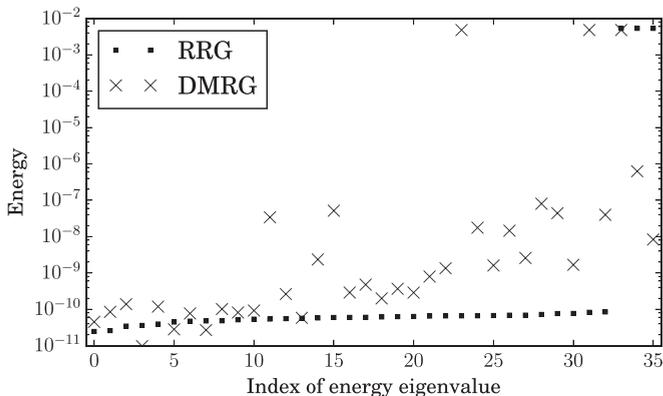


FIG. 7. Energy eigenvalues of Bravyi-Gosset model with $N = 32$ sites within the subspace obtained by RRG for $(s, D) = (36, 3)$. Also shown are DMRG results for the 36 lowest-energy states.

qubit. As the spectrum is invariant under global rotation, the Hamiltonian is fully specified by the two parameters $\theta \in [0, 2\pi)$, $p \in [0, 1/2]$. Restricting to $\theta = 0$, we may rewrite Eq. (8) in a more familiar notation:

$$H = \sum_{i=0}^{N-2} \left(\frac{\sqrt{p(1-p)}}{2} (\sigma_i^x \sigma_{i+1}^x - \sigma_i^y \sigma_{i+1}^y) + \frac{1}{4} \sigma_i^z \sigma_{i+1}^z \right) + \sum_{i=0}^{N-1} \left(\frac{1-2p}{4} \sigma_i^z + \frac{1}{8} \right). \quad (9)$$

That is, this model is equivalent to a particular XYZ model in a fine-tuned field. For any value of p the system exhibits $(N + 1)$ -fold ground-state degeneracy. Basis states for the ground space can roughly be thought of as having two regions of differing magnetization, with an interface which can be located at any site with ground-state energy $\varepsilon_0 = 0$. (Refer to Bravyi and Gosset [19] for a full description.) Therefore the algorithm choice $s \geq N + 1$ is sufficient to obtain the full ground space.

The low-energy spectrum obtained by RRG for this model at $N = 32$ is shown in Fig. 7, along with the DMRG results. We use $p = 1/2$; that is, $|\psi\rangle$ is a Bell state. Using $(s, D) = (36, 3)$, RRG identifies the full zero-energy ground space to within an accuracy determined by τ , the truncation error of the MPS. In contrast, obtaining the full ground space of this model is challenging for DMRG, which becomes hampered by candidate states of very high entanglement, often requiring a bond dimension an order of magnitude larger than those of RRG candidate states in order to achieve similar truncation error. These not only are computationally intensive to optimize, but also present DMRG with difficulty finding further excited states, as the modified Hamiltonian includes nonlocal projectors. Thus, the candidate states are not accurate eigenstates of the original Hamiltonian. This difficulty is evident in run times as well; to obtain the data shown took 10h for RRG and 40h for DMRG. Here we use DMRG without taking into account the degenerate ground-state manifold, and we consider these results to be only a point of reference. Use of a specialized approach like multiple targeting could improve accuracy, or diagonalization of the original Hamiltonian within the subspace spanned by

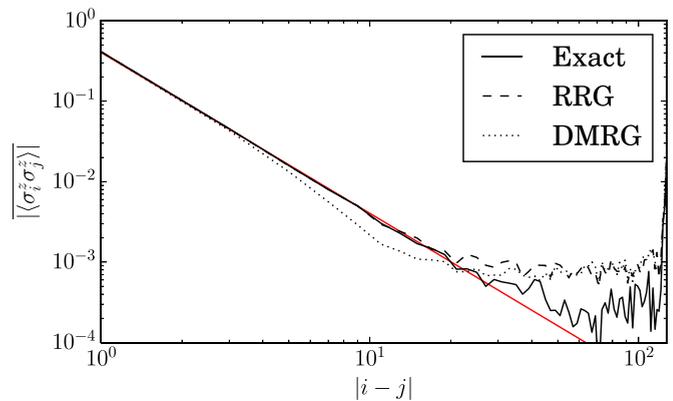


FIG. 8. Disorder-averaged decay of correlations of candidate ground states of the random XY model for $N = 128$, as compared to exact results obtained through the Jordan-Wigner transformation. The predicted power-law behavior is indicated by the red line.

the DMRG candidate states could recover much of the ground space; however, no such specialized approach is needed for RRG.

D. Random XY model

The random XY model is an inhomogeneous spin-1/2 system with Hamiltonian

$$H = \sum_{i=0}^{N-2} J_i (\sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y), \quad (10)$$

where the position-dependent coupling constants J_i are drawn from a random distribution. If the logarithm of the distribution is broad, Dasgupta-Ma real-space renormalization group analysis identifies the ground state as the random singlet phase, in which pairs of spins form singlet states at all length scales [22–25]. This model is tractable by the Jordan-Wigner transformation, which maps onto free spinless fermions. We use this system as a benchmark of algorithmic ability to encode long-range correlations in the ground state.

We use the following distribution for the Hamiltonian terms: $p(J_i) = \frac{1}{\Gamma} J_i^{-(1+1/\Gamma)}$, $J_i \in (0, 1]$, with Γ controlling the width of the distribution of log-energies [25]. We fix $\Gamma = 2$, which is sufficiently broad that the ground state is composed predominantly of localized singlet states on neighboring sites, along with spatially separated correlated qubits occurring at all length scales. As a metric we use the average two-point correlation function $\langle \sigma_i^z \sigma_j^z \rangle$ as a function of separation $r = |i - j|$ in the ground state, which is known to decay algebraically as r^{-2} . This quantity is compared to exact diagonalization results from the inhomogeneous free fermion description in Fig. 8.

These results are intended to present a fair comparison between DMRG and RRG. Both methods used unrestricted MPS bond dimension to achieve a truncation error $\tau \leq 10^{-12}$. Typically the ground-state bond dimension is similar for both methods. The RRG parameters are $(s, D) = (4, 5)$. DMRG used 20 sweeps per state, and convergence of several “hard” examples (see below) was confirmed using 50 sweeps. DMRG

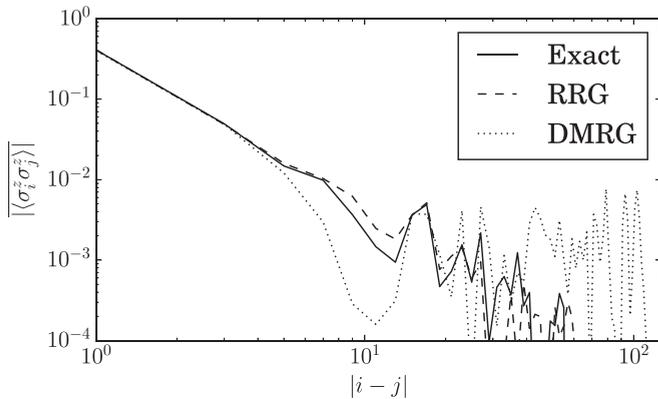


FIG. 9. A typical “hard” instance contained in the disorder average above, with energy gap $\gamma \approx 10^{-7}$. This is sufficiently large for RRG to track the long-range correlations with $(s, D) = (4, 5)$. DMRG displays a tendency for lower correlations until saturating at the noise floor.

typically took 1 h to converge $s = 4$ states and RRG took 8 h to complete. The average is over 150 disorder realizations.

The observed “saturation” of the correlations of Fig. 8 to a noise floor arises from the structure of the low-energy excited states. For a broad initial distribution, the energy gap of a specific disorder Hamiltonian may be very small. For any method using MPS, a lower limit on the gap in order to distinguish the ground state (at energy ε_0) is $\gamma \sim \tau \varepsilon_0$, below which the MPS truncation procedure will randomly select a vector from the low-lying subspace. However, even for realizations with much larger gaps a candidate ground state may include substantial contributions from low-energy excited states. A singlet of length l has energy scale $\varepsilon \sim e^{-\sqrt{l}}$; thus, the low-lying states involve excitations localized on the long-range entangled sites. Choosing a random superposition of these amounts to white noise at long distances. Instances of such Hamiltonians in the disorder average must necessarily eventually overwhelm the decay of correlations; here the distribution of energy gaps is very broad on a log scale [26], so these cases are frequent. However, in all cases the RRG candidate state has $\mathcal{O}(1)$ overlap with the true ground state, and typically this overlap is greater than 99%.

For disorder-averaged correlations at short range up to $|i - j| \approx 20$, RRG reproduces algebraic decay of correlations matching the exact results. In contrast, the DMRG candidate states demonstrate stronger decay of correlations. There is no systematic difference in MPS bond dimension between DMRG and RRG, indicating that RRG is not simply using additional resources, but is indeed more sensitive to long-range correlations.

Independent of the saturation due to the energy gap, the disorder average comprises both “easy” and “hard” instances. In easy cases both DMRG and RRG match the exact results closely at all length scales. In the hard cases both algorithms obtain the correlations only approximately, but DMRG appears to consistently underestimate correlations. RRG does not demonstrate a tendency toward either enhanced or reduced correlations. We provide an example of the spatially averaged correlations from a hard disorder realization in Fig. 9.

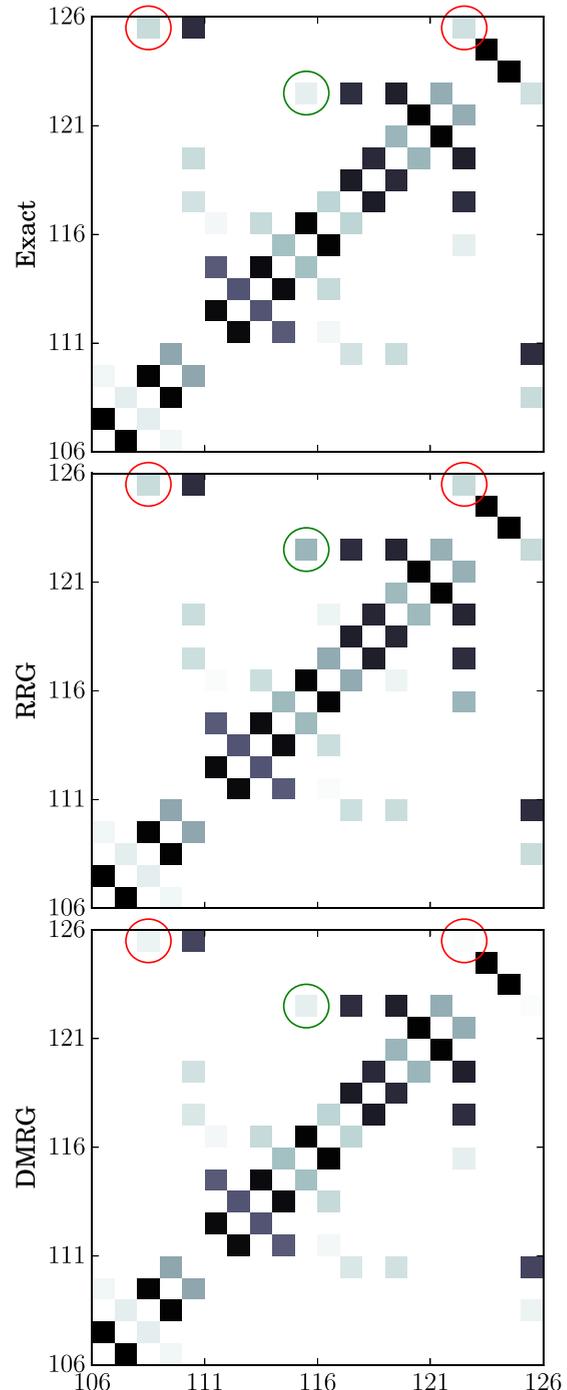


FIG. 10. Expectation value $\langle \sigma_i^z \sigma_j^z \rangle$, where sites i and j are given by the axes for $i, j \in [106, 125]$. The color scales with $\log |\langle \sigma_i^z \sigma_j^z \rangle|$ and runs from $[-2, 0]$ in all plots, with darker color indicating a higher value. The diagonal is omitted. Circles mark particular sites where differences between exact results and candidate states are evident. This disorder realization is the same “hard” instance shown in Fig. 9.

Figure 10 shows an example of measured correlations $\langle \sigma_i^z \sigma_j^z \rangle$ for various sites $i, j \in [106, 125]$ in this particular disorder realization. Each square corresponds to a measurement $\langle \sigma_i^z \sigma_j^z \rangle$ where (i, j) are specified by the axes. Darker squares indicate a larger magnitude of correlation between these sites. We show the exact results, RRG, and DMRG, and indicate some partic-

ular pairs of sites where either DMRG (red) or RRG (green) differ visibly from exact results. These variations in certain entangled sites tend toward reduced correlations in DMRG candidate ground states; it is unclear how much additional sweeping is required to compensate. RRG shows similar inaccuracies, but these are random, due to states missing from certain viable sets. Accurate correlations emerge in the disorder-averaged value, and the performance on individual disorder realizations can be controllably improved by tuning the dimension of the viable sets through the parameters s and D .

IV. DISCUSSION

DMRG has long been the method of choice for numerical calculations involving ground states of 1D systems, and over time both its efficiency and range of applicability have gone through multiple improvements and extensions. One of the main findings of our initial numerical investigation is that the RRG algorithm, developed for theoretical purposes, can in fact be made quite effective in practice, to the point of providing a potentially viable alternative to DMRG in certain cases of practical interest. We stress that the choices of parameters that we employ in our numerics are quite far from the theoretically guaranteed regime. Additionally, many of the building blocks required for the proof have been altered in our implementation. Therefore the strict guarantees no longer hold. Regardless, we find that RRG obtains ground-state candidates having large overlap with the true ground state in a variety of physically relevant models, and surpasses existing techniques in obtaining low-energy excited states and ground states of particular models demonstrating large degeneracy or long-range entanglement.

Like another numerical scheme, *time-evolving block decimation* (TEBD), the RRG procedure is a projector method, relying on operators extracted from the AGSP to guide the choice of states between scales. As a result, given a sufficiently accurate AGSP, RRG will not output a part of the spectrum strictly excited above the ground space. This is advantageous relative to variational ansatz methods which may without warning converge to an excited state rather than the ground state. (For example, if the energy landscape in Hilbert space has local minima or is very flat in the low-energy space, as is the case with the random XY model of Sec. III D.) The downsides to projector methods are that performance strongly depends on the gap and that a random initial state, even taken from the manifold of low bond dimension MPS, has exponentially small overlap with the ground state. RRG circumvents the latter issue by never choosing a trial wave function on the entire system, but rather building global states from wave functions supported on blocks which *already have good viability*; thus the projection step never has to overcome starting with an exponentially small overlap between the initial and the target state.

At present the run times required by the algorithm remain a challenge. Thus, the feasibility of RRG as a numerical method is essentially determined by the scaling discussed previously. This situation does invite future improvements. Some are immediate: for example, one may exploit symmetries of a particular problem (say, reflection symmetry across the middle of the system) in order to reduce duplication of work.

Other improvements to the current implementation are more technical. For example, as described here the management of subspaces is clumsy: operations such as addition of MPSs necessitate keeping careful track of gauge and add computational overhead for what is in principle a simple procedure. The use of data structures more appropriate to these operations could ameliorate scaling problems in all steps of the algorithm [27].

Indeed, an advantage of RRG is precisely this flexibility, to operate independently of a specific representation of states in Hilbert space. Here we have described an MPS RRG. In order to translate the logic to subspaces whose basis states are described by MERA—as would be natural for critical phases—one needs only the ability to perform evaluation of observables and addition. The former is a standard contraction which is highly efficient in MERA, and the latter can be seen as a variational process on overlaps, providing a straightforward interpretation as a MERA operation. Systems with periodic boundary conditions also present an interesting generalization, as until the final level the steps of the algorithm are insensitive to the system boundaries, provided an appropriate AGSP is given. On a more speculative note, other tensor network ansatzes may also be amenable: although it is not known that the RRG algorithm scales efficiently in higher dimensions, the hierarchical structure does generalize in an evident way and it may be the case that the algorithm gives acceptable results for projected entangled pair state representations of some two-dimensional systems.

Our numerical results suggest situations in which RRG may perform well relative to existing algorithms. The first, informed by Sec. III A, is a case in which localized and delocalized excitations lie close in energy. An optimization algorithm operating on local degrees of freedom in a sweeping pattern may exhibit a bias toward delocalized excitations, which allow for effective optimization on many lattice sites. RRG is largely insensitive to such distinctions. The second case is that of Sec. III C, exhibiting highly degenerate ground states. The full ground space is more accurately found in its entirety by RRG than DMRG. The iterative DMRG procedure of finding states is susceptible to finding poor or highly entangled candidates, which reduce the accuracy of subsequent candidates. Such a limitation is not fundamental and could likely be eliminated by modification of the procedure; however, no such modification is necessary for RRG. Finally, in Sec. III D we observe in the random XY model in the random singlet phase that long-range correlations are encoded more precisely in the ground-state candidate of RRG than of DMRG, influencing observables computed for the state.

The examples we provide illustrate specific properties indicating that a model may be well suited for RRG. However, very little is known about its more general performance: other systems with disorder, periodic boundary conditions, and higher dimensions all pose interesting challenges and could constitute exciting new directions within this formalism.

ACKNOWLEDGMENTS

We acknowledge useful discussions with M. Fishman and S. White's research group, as well as with C. White and C.-J. Lin. The numerical results were computed with the ITensor library

[20] of E. Stoudenmire and S. White. This work was supported by the Institute for Quantum Information and Matter, an NSF Physics Frontiers Center, with support of the Gordon and Betty Moore Foundation. Additional funding support was provided by the NSF through Grant No. DMR-1619696.

APPENDIX: DIFFERENCES FROM ARAD *ET AL.* [15]

In this Appendix we give a detailed account of the main points of departure of our numerical procedure from the theoretically guaranteed algorithm introduced in Arad *et al.* [15], giving heuristic justification for our choices. We refer to the paper for a more thorough introduction to the main concepts discussed here, such as the notion of viable set and AGSP.

For concreteness we base our comparison on the algorithm presented in Arad *et al.* [15] for the case of a local Hamiltonian with degenerate gapped ground space [Assumption (DG)]. The algorithm is stated as Algorithm 1 in Arad *et al.* [15]. It consists of two main steps, *Generate* and *Merge*. The two steps together recursively construct a sequence of viable sets V_m^λ for an N -qubit local Hamiltonian, where as in the main text m denotes a scale parameter and λ indexes a subregion.

1. Generate

The goal of the *Generate* step is to generate an MPO representation for a suitable AGSP. In Arad *et al.* [15] a fresh AGSP is computed for each scale m and region λ . Given a decomposition $\mathfrak{H} = \mathfrak{H}_L \otimes \mathfrak{H}_m^\lambda \otimes \mathfrak{H}_R$, a global AGSP is defined as $K_m^\lambda = T_k(\tilde{H})$, where \tilde{H} is a norm-reduced approximation of H (which depends on the region decomposition) and T_k a suitably scaled Chebyshev polynomial of degree k . The operators $A_{m,r}^\lambda$ are then computed from a specific decomposition of K_m^λ across the left and right boundaries, yielding D^2 terms $A_{m,r}^\lambda$ such that the expansion procedure $V_m^\lambda \rightarrow W_m^\lambda$ described in the main text is guaranteed to have a significant improvement on the viability parameter.

Here we depart from the theoretical algorithm in two important ways. First we use a simpler construction of AGSP, which we expect to exhibit similar behavior but is more efficient to compute. Our AGSP takes the form of an approximation $K \approx e^{-kH/t}$ obtained by Trotter decomposition. (In Arad *et al.* [15] a similar approach is taken to norm-reduce the parts of the Hamiltonian that lie in the regions L , M , and R but are a distance at least $\ell > 0$ from the boundaries.) In Arad *et al.* [15] the properties of the Chebyshev polynomial are essential to establish that the AGSP has sufficiently low bond dimension across the boundaries of region M . Considering only the efficiency in terms of improvement in viability, however, the use of $e^{-kH/t}$ over the whole chain gives similar guarantees.

Using our simpler construction implies a loss of theoretical control over the bond dimension D of the AGSP operator across the left and right cuts. This entails a second main point of departure from the theoretical algorithm, as a choice has to be made as to which operators $A_{m,r}^\lambda$ to keep. As described in the main text we proceed in a natural way by considering the MPO as an MPS and performing SVD operations to create virtual bonds between sites. We then make the choice of

keeping operators associated with the D^2 highest Schmidt weights. This choice is heuristic: the Schmidt weights control the Frobenius norm of the associated term $A_{m,r}^\lambda$, rather than the operator norm of the resulting operator, as would be desirable. The heuristic nevertheless proved effective: in practice the magnitude of the Schmidt coefficients often fell off quickly, allowing for a relatively aggressive choice of cutting point.

2. Merge process

The second step in the algorithm is called *Merge*. The goal of this step is to combine two neighboring viable sets into a single viable set over the union of the two regions, with similar approximation and size guarantees. The procedure is described as *Merge'* in Arad *et al.* [15]. *Merge'* is provided as input viable sets W_m^λ and $W_m^{\lambda+1}$ defined over neighboring regions, and returns a viable set $W_{m+1}^{\lambda/2}$ defined over the union of the two regions. *Merge* consists of three steps: *tensoring*, *random sampling*, and *error reduction*.

(1) *Tensoring*: This step is the same as in Arad *et al.* [15].

(2) *Random sampling*: Here, as already mentioned in the main text, we depart from Arad *et al.* [15] in an important way. In Arad *et al.* [15] a family of s vectors lying in $W_m^\lambda \otimes W_m^{\lambda+1}$ is obtained by random sampling within the subspace. In practice this procedure is very inefficient: (i) it requires performing high-weight (random) linear combinations of MPS, a step that is computationally expensive due to the MPS renormalization procedure; (ii) the linear combinations formed tend to be arbitrary, and in particular their MPS representations may have high MPS bond dimension, as each vector may include an ‘‘irrelevant’’ (with respect to the low-energy eigenspace of the Hamiltonian) component that artificially inflates its complexity.

Here we replace random sampling by a deterministic choice of the s lowest-energy eigenvectors of the restriction of H to $W_m^\lambda \otimes W_m^{\lambda+1}$. The idea is that low-energy eigenstates are likely, due to the local structure of the Hamiltonian, to display less entanglement. Indeed, in practice this procedure is much more efficient, and yields MPS with lower bond dimension, than the random sampling proposed in Arad *et al.* [15].

However, there is *a priori* no reason for the low-energy eigenstates of the block Hamiltonian to form a viable set for the global low-energy space. A simple heuristic argument can nevertheless be given to argue correctness of our procedure. Recall that the viability criterion Eq. (1) guarantees that the initial tensor product space supports a good approximation to any ground state. Considering the Schmidt decomposition of this approximation, each of the Schmidt vectors will have a certain energy with respect to the block Hamiltonian $H_{m+1}^{\lambda/2}$, which may not be minimal. The key is thus to argue that vectors with high energy will not have an important contribution to the Schmidt decomposition of the ground state. In general, approximation error and energy difference can scale with the norm of the Hamiltonian, making the argument difficult. However, for the purposes of approximating the ground space of a local Hamiltonian two elements play in our favor: first, locality of H , and second, the area law. The former allows one to show that the low-energy space of H is well approximated by an approximation of H with constant norm, so that the error blowup mentioned above can be controlled (see Arad

et al. [15, Proposition 3], for a precise statement). The latter establishes that the ground state has low bond dimension, so that few Schmidt vectors need to be considered (see Arad *et al.* [15, Lemma 15], for details on how this can be used). Together these two properties provide a heuristic argument in favor of our modified procedure.

(3) *Error reduction*: The goal of this step is to improve the approximation quality of the viable set. We follow

the procedure described in Arad *et al.* [15], except that the operators $\{A_{m,r}^\lambda\}$ are generated differently, as already described.

The final iteration is performed on two viable sets $V_{m^*-1}^0$ and $V_{m^*-1}^1$, each with support on one half of the system. The algorithm returns the low-lying energies and eigenstates obtained via exact diagonalization of the Hamiltonian restricted to the final viable subspace.

-
- [1] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
 [2] F. Verstraete, V. Murg, and J. Cirac, *Adv. Phys.* **57**, 143 (2008).
 [3] U. Schollwöck, *Ann. Phys. (NY)* **326**, 96 (2011).
 [4] K. G. Wilson, *Rev. Mod. Phys.* **47**, 773 (1975).
 [5] S. R. White, R. M. Noack, and D. J. Scalapino, *Phys. Rev. Lett.* **73**, 886 (1994).
 [6] X. Chen, Z.-C. Gu, and X.-G. Wen, *Phys. Rev. B* **83**, 035107 (2011).
 [7] M. Levin and C. P. Nave, *Phys. Rev. Lett.* **99**, 120601 (2007).
 [8] Z.-C. Gu, M. Levin, and X.-G. Wen, *Phys. Rev. B* **78**, 205116 (2008).
 [9] G. Evenbly and G. Vidal, *Phys. Rev. Lett.* **115**, 180405 (2015).
 [10] G. Vidal, *Phys. Rev. Lett.* **99**, 220405 (2007).
 [11] G. Evenbly and G. Vidal, *Phys. Rev. B* **79**, 144108 (2009).
 [12] R. N. C. Pfeifer, G. Evenbly, and G. Vidal, *Phys. Rev. A* **79**, 040301 (2009).
 [13] G. Evenbly and G. Vidal, *Phys. Rev. B* **81**, 235102 (2010).
 [14] G. Evenbly and G. Vidal, *Phys. Rev. Lett.* **115**, 200401 (2015).
 [15] I. Arad, Z. Landau, U. Vazirani, and T. Vidick, *Commun. Math. Phys.* **356**, 65 (2017).
 [16] Z. Landau, U. Vazirani, and T. Vidick, *Nat. Phys.* **11**, 566 (2015).
 [17] I. Arad, Z. Landau, and U. Vazirani, *Phys. Rev. B* **85**, 195145 (2012).
 [18] M. Suzuki, *Prog. Theor. Phys.* **56**, 1454 (1976).
 [19] S. Bravyi and D. Gosset, *J. Math. Phys.* **56**, 061902 (2015).
 [20] E. M. Stoudenmire and S. R. White, ITensor — Intelligent Tensor Library, <http://itensor.org>.
 [21] C.-J. Lin and O. I. Motrunich, *Phys. Rev. A* **95**, 023621 (2017).
 [22] R. N. Bhatt and P. A. Lee, *Phys. Rev. Lett.* **48**, 344 (1982).
 [23] S.-k. Ma, C. Dasgupta, and C.-k. Hu, *Phys. Rev. Lett.* **43**, 1434 (1979).
 [24] C. Dasgupta and S.-k. Ma, *Phys. Rev. B* **22**, 1305 (1980).
 [25] D. S. Fisher, *Phys. Rev. B* **50**, 3799 (1994).
 [26] D. S. Fisher and A. P. Young, *Phys. Rev. B* **58**, 9131 (1998).
 [27] Recently, the authors have significantly advanced the state of the RRG code, which may be found online at <https://github.com/brendenroberts/RigorousRG>.