

Role of conservation laws in the density matrix renormalization group

Thomas G. Kiely^{✉*} and Erich J. Mueller[†]

Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, New York 14853, USA



(Received 7 July 2022; revised 20 October 2022; accepted 6 December 2022; published 14 December 2022)

We explore matrix product state approximations to wave functions which have *spontaneously broken* symmetries or are *critical*. We are motivated by the fact that symmetries, and their associated conservation laws, lead to block-sparse matrix product states. Numerical calculations which take advantage of these symmetries run faster and require less memory. However, in symmetry-broken and critical phases the block-sparse *ansatz* yields less accurate energies. We characterize the role of conservation laws in matrix product states and determine when it is beneficial to make use of them.

DOI: [10.1103/PhysRevB.106.235126](https://doi.org/10.1103/PhysRevB.106.235126)

I. INTRODUCTION

Our most powerful numerical techniques for studying one-dimensional (1D) quantum systems, such as the density matrix renormalization group (DMRG) or its time-dependent generalizations [1,2], are based upon a systematic truncation of the entanglement between neighboring regions of space. Within these approaches, the quantum-mechanical wave function has the structure of a matrix product state (MPS): the amplitude of any given configuration is calculated by taking the product of a series of matrices, one for each site. In the presence of symmetries, these matrices can be taken to be *block sparse*, where the majority of matrix elements vanish. This structure is used in all modern codes to accelerate performance. Here we assess the limitations of this block-sparse structure: What happens when the symmetry is spontaneously broken, or if one is at a critical point? Using the transverse-field Ising model as a pedagogical example, we elucidate how the most efficient description of a state in the symmetry-broken phase does not make use of conservation laws. We then explore critical systems: In the superfluid phase of the 1D Bose-Hubbard model and the metallic phase of the 1D Fermi Hubbard model, we find that an MPS which respects the symmetry requires larger matrices to achieve the same accuracy. For some parameter ranges, this results in a larger memory footprint and longer runtime.

According to Noether's theorem, symmetries are closely related to conservation laws [3,4]. As a relevant example, consider a Hamiltonian that is invariant under the transformation

$$\hat{H} \rightarrow \hat{U} \hat{H} \hat{U}^\dagger, \quad (1)$$

where $\hat{U}(\theta) = e^{i\theta \hat{N}}$ and $\hat{N}$ is the total particle-number operator. Equation (1) defines a continuous $U(1)$ symmetry, parametrized by θ . This can only be true if $[\hat{H}, \hat{N}] = 0$, which

is the formal quantum-mechanical statement that $\hat{N}$ is conserved. Consequently, we can find simultaneous eigenstates of $\hat{H}$ and $\hat{N}$.

One of the most profound features of many-body physics is that, in the thermodynamic limit, the symmetry may be *spontaneously broken* [5]: An infinitesimal symmetry-breaking field leads to a ground state which is neither invariant under the symmetry operation, nor is it an eigenstate of the conserved charge. A relevant example is a Bose-Einstein condensate, which chooses a particular phase and contains an indefinite number of particles. Spontaneous symmetry breaking is always associated with a ground-state degeneracy, and one can restore the symmetry by taking an appropriate quantum superposition of the degenerate ground states. In the case of a discrete symmetry, such symmetry-restored states are “Schrödinger cats.”

In Sec. III A we present the transverse-field Ising model, which possesses a discrete $\mathbb{Z}_2$ symmetry, as a pedagogical example. It has two zero-temperature phases: a paramagnetic phase, in which the ground state respects the symmetry, and a ferromagnetic phase, which breaks it. In both phases one can use a symmetry-preserving MPS to describe the ground state. In the ferromagnetic phase, however, the resulting MPS corresponds to the aforementioned Schrödinger cat, which is a superposition of the two symmetry-broken solutions. These symmetry-broken constituents are less entangled than the symmetry-preserving Schrödinger cat, and hence are more efficient to express as a MPS [1]. Aspects of this behavior are known by the community, but rarely discussed in the literature.

The situation is far more complicated for continuous symmetries. One-dimensional systems with short-ranged interactions and finite susceptibilities cannot break a continuous symmetry [6–8]. Instead, strong quantum fluctuations lead to correlation functions that fall off as a power law [9]. It is far from obvious if they are better described by an MPS that respects the symmetry or one that explicitly breaks it. We consider two examples: the Bose-Hubbard model and the Fermi-Hubbard model. In both cases, we find that the symmetry-conserving MPS requires a larger bond dimension

*tgk37@cornell.edu

†em256@cornell.edu

(the linear size of the MPS matrices) to achieve the same accuracy as the symmetry-broken MPS. For the Fermi-Hubbard case the improvement is fairly modest, while for the Bose-Hubbard case it is quite substantial. We show that the key difference is the scaling of density fluctuations, which is characterized by the Luttinger parameter K , and we quantify this relationship.

Throughout this paper we largely compare sparse and dense MPS representations of states with the *same bond dimension*. This allows us to cleanly understand the ways in which conservation laws manifest in critical and symmetry-broken states. We emphasize that the sparse MPS states with bond dimension χ are a subset of the dense states with the same bond dimension. Thus, imposing the symmetry can never improve the energy of the variational ground state. The important question is the *extent* to which the dense MPS is a better variational *ansatz*.

For practical numerical calculations, one is likely interested in a different question: For a fixed numerical accuracy, does it take more computer time to calculate the ground state using a sparse or dense MPS? One could similarly ask about memory or disk usage. Unfortunately, these questions will inevitably depend on details of the implementation: How does one store the block-sparse matrices? How does one implement basic linear algebra operations? We largely relegate these practical questions to Appendix A. We find that within the ITENSOR package [10], runtimes can be either increased or reduced by imposing conservation laws, depending on parameters. One would intuitively guess that proximity to symmetry breaking would determine the extent to which one benefits from the block-sparse structure. For the systems we study, the Luttinger parameter K quantifies this proximity: the ideal Bose gas, with $K = 0$, has off-diagonal long-range order and is often interpreted as a symmetry-broken state [11]. As expected from this argument, the benefits from using dense tensors are largest at small K . The other relevant parameter is the target accuracy, which determines the bond dimension. For low accuracy (small χ), dense calculations are faster, while for high accuracy, block-sparse calculations are faster. The crossover point depends on K : smaller K favors the dense *ansatz*.

In Sec. II we review features of the MPS *ansatz* and discuss how Abelian symmetries lead to block-sparse MPS tensors. In Sec. III we present our results: we begin with the transverse-field Ising model as a pedagogical example (Sec. III A), then move to the more nuanced Bose-Hubbard (Sec. III B) and Fermi-Hubbard (Sec. III C) models. In Sec. IV we present a more general interpretation of the results in Secs. III B and III C in terms of the Luttinger parameter. We conclude in Sec. V.

II. CONSERVATION LAWS IN MPS

An MPS incorporates conservation laws by placing restrictions on which matrix elements can be nonzero [12]. This sparse structure can be exploited to dramatically speed up tensor contractions. For the purpose of this paper, we will only consider Abelian symmetries generated by a global operator $\hat{Q} = \sum_i \hat{Q}_i$ that commutes with the Hamiltonian: $[\hat{H}, \hat{Q}] = 0$. Here $\hat{Q}_i$ are a set of mutually commuting single-site opera-

tors, where i indexes the sites of the MPS in real space: for concreteness, one can envision $\hat{Q} = \hat{N}$, as in Eq. (1), and take $\hat{Q}_i = \hat{N}_i$ to be the number of particles on site i .

We consider a matrix product state wave function on L sites. The MPS *ansatz* can be schematically written as $|\psi\rangle = \sum_{\sigma} A^{\sigma_1} A^{\sigma_2} \dots A^{\sigma_L} |\sigma_1, \sigma_2, \dots, \sigma_L\rangle$ where A^{σ_i} corresponds to a matrix with elements $(A_i)_{s_{i-1}s_i}^{\sigma_i}$. The sum is taken over all $\sigma_1, \sigma_2, \dots, \sigma_L$, where σ_i corresponds to the allowed states on site i , and over shared indices s_i between adjacent matrices. Here s_{i-1} and s_i are the left and right MPS bond indices. The bond dimension χ is the number of different possible values of s . In order to make use of the conservation law we write the local Hilbert space in the eigenbasis of the local operator $\hat{Q}_i$, and define a function $q(\sigma_i)$ which associates a *charge* with each of the local basis states: $\hat{Q}_i |\sigma_i\rangle = q(\sigma_i) |\sigma_i\rangle$. We similarly associate a charge with each possible value of the bond indices. The conservation law is imposed by requiring that the only nonzero elements of A_i obey

$$q(s_{i-1}) + q(\sigma_i) - q(s_i) = 0. \quad (2)$$

In the case of number conservation, one can interpret $q(s_{i-1})$ as the number of particles to the left of site i , and $q(s_i)$ as the number to the left of site $i + 1$. For infinite chains, it is convenient to define $q(\sigma_i)$ as the deviation of the quantum number from its average so that the charges of the bond indices are more readily truncated.

As should be clear, the block-sparse condition in Eq. (2) greatly reduces the number of matrix elements which need to be stored and speeds up all matrix operations. Its limitations, however, are illustrated by considering a simple Gutzwiller mean-field wave function: $|\Psi\rangle = \prod_{i=1}^L \otimes (a|0\rangle_i + b|1\rangle_i)$, which represents a Bose-Einstein condensate in which each site contains the superposition of 0 and 1 particle. This is a MPS with bond dimension $\chi = 1$, but it does not obey Eq. (2). One can rewrite it using the conservation laws, but that comes at the cost of greatly increasing χ . For a chain of length L , for example, one needs $\chi = L$, and the MPS matrices can take the form

$$A_i = \begin{pmatrix} a|0\rangle_i & b|1\rangle_i & & & \\ & a|0\rangle_i & b|1\rangle_i & & \\ & & a|0\rangle_i & b|1\rangle_i & \\ & & & \ddots & \ddots \end{pmatrix}. \quad (3)$$

The rows correspond to configurations where there are 0, 1, 2, ... particles to the left of this site. The bond index increments whenever a site is occupied.

III. RESULTS

We characterize the distinction between a quantum-number-conserving (sparse) MPS and a nonconserving (dense) MPS by running iDMRG simulations on a few well-known models. We begin with a pedagogical discussion of the transverse-field Ising in Sec. III A, which exhibits discrete spontaneous symmetry breaking. In this particular model, which has been studied extensively with a wide range of analytical and numerical techniques [13–16], we show how one can explicitly construct the dense MPS out of the sparse MPS (and vice versa). This transformation preserves the variational energy but not the bond dimension, and hence yields insight

into the relative efficiency of the dense and sparse *ansätze*. We then move to examples of Luttinger liquids, namely, the Bose-Hubbard (Sec. III B) and Fermi-Hubbard (Sec. III C) models. These are more complicated systems that do not explicitly break any symmetries, so they necessitate more detailed numerical comparisons. We make use of the ITENSOR library for an efficient implementation of quantum number conservation [10]. Further details of the numerical simulations are discussed in Appendix B.

A. Transverse-field Ising model: A pedagogical example

The one-dimensional transverse-field Ising model is an exactly solvable model of $s = \frac{1}{2}$ spins on a lattice with nearest-neighbor interactions. The Hamiltonian is given by

$$\mathcal{H}_{\text{TFI}} = -J \sum_j (\sigma_j^x \sigma_{j+1}^x + \alpha \sigma_j^z), \quad (4)$$

where σ_j^β is the Pauli spin matrix ($\beta = x, y, z$) acting on the spin on site j . The ratio of the transverse field strength to the nearest-neighbor interaction strength α is the only nontrivial parameter in the ground-state phase diagram (here we consider the ferromagnetic model: $J, \alpha > 0$). While $\mathcal{H}_{\text{TFI}}$ does not conserve total magnetization, it has a global $\mathbb{Z}_2$ symmetry [$\mathcal{H}_{\text{TFI}}, \hat{P}$] = 0, where the parity operator $\hat{P} = \prod_j \sigma_j^z = \exp[i\pi \sum_j (\sigma_j^z - 1)/2]$, rotates all spins about the $\hat{z}$ axis by π . This parity symmetry implies that the magnetization along the $\hat{z}$ direction is conserved modulo 2.

The transverse-field Ising model has two zero-temperature phases. When $\alpha > 1$, the σ^z term dominates and spins tend to align with the transverse field. This phase is even under parity transformations: $\hat{P}|\psi_{\alpha>1}\rangle = |\psi_{\alpha>1}\rangle$. When $\alpha < 1$, the exchange term dominates and spins will tend to align with one another in the $\pm\hat{x}$ direction. In the thermodynamic limit, an infinitesimal field in the $\hat{x}$ direction will result in a ground state $|+\rangle$ with a finite magnetization in the $+\hat{x}$ direction. This is an example of spontaneous symmetry breaking: the state $|-\rangle = \hat{P}|+\rangle$ is orthogonal to $|+\rangle$ and has a magnetization in the $-\hat{x}$ direction. Two parity-conserving ground states can be formed by taking $|+\rangle \pm |-\rangle$. In the limit $\alpha \rightarrow 0$, parity-conserving ground states are given by Greenberger-Horne-Zeilinger (GHZ) states [17],

$$\frac{1}{\sqrt{2}}(|\rightarrow\rightarrow\rightarrow\ldots\rangle \pm |\leftarrow\leftarrow\leftarrow\ldots\rangle), \quad (5)$$

where $|\rightarrow\rangle_j$ denotes a spin on site j oriented in the $\hat{x}$ direction.

We use the infinite DMRG algorithm (iDMRG) to find the ground state of Eq. (4) as a function of α . We separately run the algorithm with and without enforcing parity conservation, using appropriate initial conditions. In Fig. 1(a) we plot the resulting entanglement entropy across a bipartition of the infinite chain. Filled red dots denote the behavior of the parity-conserving MPS, while open blue circles show the non-parity-conserving results. For $\alpha > 1$ the *ansätze* converge to the same state, which has zero entanglement entropy as $\alpha \rightarrow \infty$. This parent state is simply a product state with all spins oriented in the $\hat{z}$ direction. The entanglement entropy diverges at the critical point $\alpha = 1$. For $\alpha < 1$, the two simulations converge to distinct, degenerate ground states. As shown




FIG. 1. (a) Entanglement entropy of the parity-conserving and parity-non-conserving MPS as a function of the dimensionless transverse field α . The model has a quantum critical point at $\alpha = 1$ where the entanglement entropy diverges logarithmically. For $\alpha > 1$ the ground state has a definite parity and both *ansätze* agree with one another. For $\alpha < 1$, the model is in the symmetry-broken phase. The parity-conserving *ansatz* must remain in an even-parity state, and thus converges to a GHZ state. The parity-non-conserving state is able to break the $\mathbb{Z}_2$ symmetry and converge to a lower-entropy state. As shown in the figure, the difference in entropy for $\alpha < 1$ is precisely $\ln(2)$. (b) Singular values of both *ansätze* at $\alpha = 0.9$. The values mirror one another, but the parity-conserving *ansatz* has exactly double the number of singular values. (c) Singular values of both *ansätze* at $1/\alpha = 0.9$. Here both *ansätze* converge to the same definite-parity state, and hence their singular values are identical.

by the arrows, the parity-conserving *ansatz* has exactly $\ln(2)$ more entanglement entropy at every point with $\alpha < 1$. This relationship is expected when the parity-conserved state is a simple superposition of the two symmetry-broken ground states. We note that this feature of unconstrained DMRG, in which the algorithm converges to the minimally entangled degenerate ground state, is generic and has been recognized in the context of topological systems [18,19].

We investigate this correspondence more closely in Figs. 1(b) and 1(c), where we plot the spectrum of singular values λ_i at representative points in both phases: λ_i^2 are the eigenvalues of the reduced density matrix when one traces over half the chain. Figure 1(b) shows the spectrum at a representative point in the symmetry-broken phase. The spectrum is effectively doubled by conserving parity: each blue singular value matches up with exactly two red singular values in each degenerate plateau. This is precisely what one would expect by taking a superposition of symmetry-broken states.

Note that the slight offset between corresponding blue and red plateaus is due to the normalization condition $\sum_i \lambda_i^2 = 1$, and the fact that the parity-conserving *ansatz* has twice as many singular values. By contrast, Fig. 1(c) shows that the singular values of both states match up perfectly when $\alpha > 1$.

A consequence of this spectral doubling is that the symmetry-preserving MPS in the symmetry-broken phase needs twice the bond dimension to yield the same accuracy as the wave function which explicitly breaks the symmetry. The MPS tensors therefore contain four times as many matrix elements, only *half* of which are eliminated by the block-sparseness condition in Eq. (2). Thus, instead of making the calculation more efficient, enforcing parity conservation requires storing twice as many matrix elements. On the paramagnetic side of the transition, the situation reverses, and the parity-conserving *ansatz* requires half as many elements.

The transverse-field Ising model is simple, and the algorithmic costs and benefits here are small. Nonetheless, it provides a clear illustration of how spontaneous symmetry breaking interacts with conservation laws in DMRG.

B. Bose-Hubbard model

The Bose-Hubbard model is a paradigmatic strongly interacting model of lattice bosons. In one dimension (1D) the Hamiltonian is

$$\mathcal{H}_{\text{BH}} = -t \sum_j (a_j^\dagger a_{j+1} + \text{H.c.}) + U \sum_j n_j n_j, \quad (6)$$

where $a_j^{(\dagger)}$ is a bosonic annihilation (creation) operator on site j of a lattice, and $n_j = a_j^\dagger a_j$ is the number operator. We focus on the superfluid phase, which in 1D is a critical phase described by Tomonaga–Luttinger-liquid theory [9, 20–22]. Unlike a Bose-Einstein condensate, it does not spontaneously break $U(1)$ gauge invariance. There is, however, quasi-long-range order corresponding to a power-law decay of the single-particle density matrix. In contrast to the phases in Sec. III A, it is not *a priori* obvious whether this superfluid phase would be better described by a variational wave function that breaks or conserves particle-number conservation.

The most interesting part of the phase diagram is near the Berezinskii-Kosterlitz-Thouless (BKT) transition at the tip of the Mott lobe. Thus, we focus on the point $U/t = 3$ with an average of $\bar{n} = 1$ particles per site. We use the standard iDMRG algorithm. For our particle-number-conserving simulations, fixing the average density is trivial, while in our unrestricted simulations we add an extra step in each iteration which corrects the chemical potential μ . This procedure is described in Appendix B 2. As described in Ref. [22], this gapless, critical phase is best analyzed using “finite entanglement scaling,” meaning that one understands the properties of the state by considering a sequence of bond dimensions χ .

Our results are summarized in Fig. 2. In the main panel we plot the density matrix $\langle a_i^\dagger a_j \rangle$ as a function of spatial separation $|i - j|$. For the number-conserving *ansatz*, the correlation function falls off exponentially at sufficiently long distances. For the dense *ansatz*, the correlation function instead approaches a constant. This constant corresponds to a Bose-Einstein condensate, indicating that the infinite bond dimension approximation spontaneously breaks the symmetry




FIG. 2. Density matrix of the 1D Bose-Hubbard model $\langle a_i^\dagger a_j \rangle$ versus spatial separation $|i - j|$ on a log-log scale for a variety of bond dimensions χ , denoted by color. Model parameters are $U/t = 3$, $\bar{n} = 1$. Curves utilizing the block-sparse particle-number-conserving MPS bend downward, while the unrestricted dense MPS curves exhibit a plateau corresponding to Bose-Einstein condensation. Dashed line shows the asymptotic power-law decay based on a scaling analysis (see Appendix C). Inset: Variational energy in units of U versus bond dimension on a log-log scale for dense and sparse matrix product states. The dense MPS achieves a lower energy for all bond dimensions, while both curves exhibit scaling consistent with Eq. (8), whose slope is given by the thick gray line.

even though exact ground state is critical. One typically refers to this phenomenon as “quasicondensation,” characterized by a quasicondensate density ρ_{qc} . We discuss this at greater length in Sec. IV. The dashed black line shows the asymptotic power-law scaling of the density matrix

$$\langle a_i^\dagger a_j \rangle \propto |i - j|^{-K/2}, \quad (7)$$

where we used a scaling analysis to find the Luttinger parameter K (see Appendix C). The dense MPS better captures the correlations: while both curves eventually bend away from the dashed line, the dense MPS curves show approximate power-law decay out to distances almost an order of magnitude larger than those of the sparse MPS.

The inset of Fig. 2 shows the variational energy of the dense and sparse *ansätze* as a function of bond dimension on a log-log scale. The dense MPS has a substantially lower energy at each bond dimension, while both curves exhibit power-law scaling of the form

$$E(\chi) = E_0 + A/\chi^{2\kappa} + \dots, \quad A > 0 \quad (8)$$

where E_0 is the true ground-state energy in the thermodynamic limit. As argued in Ref. [23], one expects that $\kappa = 6/(c + \sqrt{12}c)$ for an MPS approximation of a conformal critical point with central charge c . The low-energy description of the Bose-Hubbard model takes the form of a single-component

Luttinger liquid, which is a conformal field theory with $c = 1$. This theoretical prediction is given by the shaded gray line in the inset, clearly showing that the data are in close agreement with Eq. (8). The dense *ansatz* is roughly an order of magnitude more accurate for the same bond dimension.

In contrast to the energy, the correlation length behaves counterintuitively. We define

$$\xi^2 = \frac{\sum_j j^2 (\langle a_j^\dagger a_0 \rangle - \langle a_j^\dagger \rangle \langle a_0 \rangle)}{\sum_j \langle a_j^\dagger a_0 \rangle - \langle a_j^\dagger \rangle \langle a_0 \rangle}, \quad (9)$$

which is the characteristic length scale of the fluctuations. Despite the fact that the number-nonconserving *ansatz* yields a density matrix which is closer to the exact result (which has an infinite correlation length), its correlation length is shorter. This unexpected result is a consequence of subtracting off the constant term in Eq. (9).

In addition to the U(1) symmetry described here, at the BKT point the Bose-Hubbard model exhibits an emergent $\mathbb{Z}_2$ particle-hole symmetry. As the number-conserving *ansatz* encodes particle and hole fluctuations with different singular values, this implies that near the BKT point many of its singular values will have nearly degenerate partners. Similar degeneracies have been used to detect forms of order [24–26]. They also indicate that the *ansatz* contains redundant information. These degeneracies do not show up in the singular values for the nonconserving *ansatz*.

C. Fermi-Hubbard model

The 1D Fermi-Hubbard model describes spin- $\frac{1}{2}$ lattice fermions with onsite interactions and Hamiltonian

$$\mathcal{H}_{\text{FH}} = -t \sum_{j,\sigma} (c_{j,\sigma}^\dagger c_{j+1,\sigma} + \text{H.c.}) + U \sum_j n_{j,\uparrow} n_{j,\downarrow}. \quad (10)$$

Here $c_{j,\sigma}^{(\dagger)}$ is a fermionic annihilation (creation) operator for a particle with spin σ on site j and $n_{j,\sigma} = c_{j,\sigma}^\dagger c_{j,\sigma}$ is the number operator. Like the 1D Bose-Hubbard model, the ground state of the 1D Fermi-Hubbard model is either a Mott insulator or a Luttinger liquid. We will again focus on the latter phase. This model is exactly solvable via the Bethe *ansatz* [27,28].

At half-filling (one particle per site) this model is in the Mott insulator phase for any $U/t > 0$. Thus, we work at quarter filling and zero net magnetization, $\bar{n}_\uparrow = \bar{n}_\downarrow = \frac{1}{4}$. Our block-sparse simulations conserve both the total particle number and the total magnetization. Number-conserving simulations at a fractional filling p/q , where p and q are integers, requires a unit cell of length qm sites where $m \in \mathbb{Z}^+$.¹ The dense MPS simulations have no restriction on the allowed unit-cell size. For the purpose of providing a reliable comparison between methods, we perform both the number-conserving and non-number-conserving simulations with a

¹It is in fact possible to work with smaller unit-cell sizes, but this will effectively encode the larger unit cell by cycling through different blocks of the MPS. This enlarges the bond dimension artificially without improving accuracy, which is detrimental to algorithmic efficiency.




FIG. 3. (a) Momentum distribution function for $\uparrow$ spins, $\langle c_{k\uparrow}^\dagger c_{k\uparrow} \rangle$, in the vicinity of $k_F = \pi/4$ for the 1D Fermi-Hubbard model. Model parameters are $U/t = 4$, $\bar{n}_\uparrow = \bar{n}_\downarrow = \frac{1}{4}$. Top and bottom plots show data from the dense and sparse *ansatz*. All curves display the expected power-law scaling [see Eq. (11)] up to broadening of the power-law singularity due to the finite MPS correlation length. The distribution function is considerably sharper for the dense MPS, indicating a longer correlation length. (b) Derivative of the momentum distribution function $\partial \langle n_{k\uparrow} \rangle / \partial k$ versus k . The axes have been rescaled by $\chi^{\kappa(\alpha-1)}$ and χ^κ , respectively, where κ and α are defined in the main text. The dense and sparse curves from (a) exhibit distinct scaling collapses, as shown. (c) Variational energy in units of U versus bond dimension on a log-log scale for dense and sparse matrix product states. As with the Bose-Hubbard model, we see the dense MPS provides a lower variational energy and that the energies of both *ansätze* are described by a power law consistent with Eq. (8) (gray line).

unit cell of four sites. A good discussion of multisite iDMRG can be found in Ref. [29].

Figure 3(a) shows the momentum distribution function for up spins $\langle n_{k\uparrow} \rangle$ in the vicinity of $k_F = \pi/4$. Both data sets show a step at k_F that grows increasingly sharp with bond dimension. Note that on this scale the distribution never goes to 0 or 1. The step height is somewhat analogous to the Fermi-liquid quasiparticle weight Z . As $\chi \rightarrow \infty$ the distribution function near $k = k_F$ should approach [9]

$$\langle n_{k,\sigma} \rangle \approx \frac{1}{2} - \text{sign}(k - k_F) |k - k_F|^\alpha, \quad (11)$$

where the exponent in the power-law singularity depends on the Luttinger parameter for charge degrees of freedom K_ρ :

$$\alpha = (K_\rho + 1/K_\rho - 2)/4. \quad (12)$$

For $U/t = 4$ and $\bar{n}_\uparrow = \bar{n}_\downarrow = \frac{1}{4}$ the Bethe *ansatz* solution gives $K_\rho \approx 1.4$ [28].

At $k = k_F$, the derivative $\frac{\partial \langle n_{k\uparrow} \rangle}{\partial k}$ in Eq. (11) diverges. For finite χ , this singularity is cut off and one instead expects $\frac{\partial \langle n_{k\uparrow} \rangle}{\partial k} \sim \chi^{\kappa(1-\alpha)}$ where $\kappa = 3/(1 + \sqrt{6})$ is the conformal scal-

ing exponent corresponding to a central charge $c = 2$ [23]. The width of the deviation from Eq. (11) scales as $\delta q \propto \chi^{-\kappa}$. Figure 3(b) demonstrates the resulting scaling collapse: For a given *ansatz*, all of the curves from Fig. 3(a) lie on top of one another. We use the theoretical values of κ and K_ρ , without any free parameters.

Strikingly, in Fig. 3(b), the dense and sparse MPS exhibit two distinct scaling collapses, the former notably sharper than the latter. Thus, while both data sets exhibit the expected conformal scaling, the dense MPS yields wave functions with a sharper singularity at k_F . These two scaling collapses can be made to line up with one another by rescaling the bond dimension by a factor of 1.8.

In Fig. 3(c) we plot the variational energy as a function of the bond dimension on a log-log scale. We again find that the energy obtained by the dense MPS is lower than that of the sparse MPS. Fixing the bond dimension, the ratio of the errors in the energy for the two *ansätze* is roughly 1.5. The shaded gray line denotes the scaling behavior in Eq. (8), which is clearly consistent with both data sets.

IV. DISCUSSION

In the symmetry-broken phase of the transverse-field Ising model, the twofold-degenerate ground-state manifold is spanned by symmetry-broken states $|\pm\rangle$, or symmetry-preserving states $|+\rangle \pm |-\rangle$. The symmetry-broken states have a smaller entanglement entropy, and hence can be described by a MPS with smaller bond dimension. The matrices in the symmetry-preserving MPS, however, are sparse.

The situation is more complicated in Secs. III B and III C, where we explored critical Luttinger-liquid states. In the thermodynamic limit these have infinite entanglement entropy, and hence an exact representation would require an MPS with infinite bond dimension. For finite bond dimension, the dense MPS *ansatz* breaks the U(1) gauge symmetry and exhibits quasicondensation. Similar to the Ising model example, one can construct a number-conserving state with density $\bar{n}$ by averaging over all values of the broken symmetry: $|\bar{n}\rangle = \int d\theta \exp[i\theta(\hat{n} - \bar{n})]|\psi_0\rangle$. Unfortunately, the $|\bar{n}\rangle$ constructed in this manner will have infinite bond dimension. This points towards a more complex relationship between the number-conserving and symmetry-broken MPS approximants. Nonetheless, for a fixed bond dimension, the symmetry-broken wave function yields a more accurate energy. As with the case of the transverse-field Ising model, this increase in accuracy comes with the cost of requiring the use of dense matrices.

The symmetry breaking found at finite χ is analogous to the quasicondensation seen in 1D Bose gases confined in traps of length L [30]. Matrix product states with finite bond dimension always have a finite correlation length ξ , and this length scale plays a similar role to L . Just as our quasicondensate density vanishes as $\chi \rightarrow \infty$, these physical systems have ρ_{qc} vanish as $L \rightarrow \infty$.

In the Fermi-Hubbard model, the quasicondensation discussed above corresponds to fictitious bosons which are constructed via a Jordan-Wigner transformation. This therefore corresponds to a topological order in the fermionic system, which is revealed via a string correlation function.




FIG. 4. Ratio of leading coefficients in the energy scaling function [Eq. (8)], $\Gamma = A_{QN}/A_{\text{dense}}$, versus the Luttinger parameter. Blue, orange, and green data correspond to the Bose-Hubbard model at particle densities $\bar{n} = 1$, $\frac{1}{2}$, and $\frac{1}{4}$, respectively. Strikingly, these data lie on a single scaling function, independent of the microscopic parameters. The function exhibits a power-law divergence as $K \rightarrow 0$, where the ground state is proximate to a Bose-Einstein condensate, while it should approach 1 as $K \rightarrow \infty$.

There is no obvious way to experimentally measure this topological quasiorder.

Comparing the inset of Fig. 2 with Fig. 3(c), it is clear that the advantage gained from breaking the symmetry is larger for the bosons than for the fermions. As noted in Eq. (8), the leading deviation of the variational energy is $\delta E = A\chi^{-2\kappa}$. Here A is smaller for the dense *ansatz*, and the improvement in accuracy from using the dense *ansatz* is quantified by the dimensionless ratio $\Gamma = A_{QN}/A_{\text{dense}}$. To achieve a fixed error in the energy, the number-conserving *ansatz* requires a bond dimension which is $\Gamma^{1/2\kappa}$ times larger than the dense *ansatz*. In our bosonic example (Sec. III B), $\Gamma = 6.7$, while in the fermionic one (Sec. III C), $\Gamma = 1.5$.

The reason for this difference is that the fermionic system has much smaller density fluctuations. In the number-conserving *ansatz*, there is a configurational entropy associated with number fluctuations between two halves of the system, requiring a larger bond dimension. The scale of these number fluctuations is set by the Luttinger parameter: A region of size L will have fluctuations $\langle(\hat{N} - \langle\hat{N}\rangle)^2\rangle \sim K^{-1} \ln L$ [31]. In a MPS of fixed bond dimension, the correlation length ξ plays the role of L . In our examples $K^{\text{Bose}} = 0.42$ is much smaller than $K_\rho^{\text{Fermi}} = 1.4$.

In Fig. 4 we show how Γ depends on Luttinger parameter for the Bose-Hubbard model at three different fillings: $\bar{n} = 1$, $\frac{1}{2}$, and $\frac{1}{4}$. We find that these data collapse onto a single universal curve which diverges as a power law $\Gamma \propto K^{-1.27(2)}$ for small K . The maximum value of K in the superfluid phase of the Bose-Hubbard model is $K = 1$, beyond which the system undergoes a Mott transition. If one were to continue the scaling function out to $K > 1$, e.g., with the inclusion of long-range interactions, one would expect the power-law behavior to break down so that $A_{\text{dense}}/A_{QN} \rightarrow 1$ as $K \rightarrow \infty$.

The metallic phase of the 1D Fermi-Hubbard model with $U > 0$ has distinct Luttinger parameters for spin (K_σ) and charge (K_ρ) degrees of freedom. The spin Luttinger parameter is fixed at $K_\sigma = 1$, while the charge Luttinger parameter $K_\rho > 1$. Both spin and density fluctuations are relevant here, so the fermionic results do not collapse onto the bosonic data in Fig. 4.

V. CONCLUSIONS

Conservation laws allow one to write matrix product states in a block-sparse manner [Eq. (2)]. It is not, however, always favorable to take advantage of this structure. For example, if the ground state spontaneously breaks the symmetry, then the resulting MPS contains redundant information whose only purpose is to impose the constraint.

These considerations are particularly interesting for critical Luttinger liquid phases, where the symmetry is *almost* broken. We find that for fixed bond dimension one more accurately estimates the ground-state energy by using a dense *ansatz* that does not rely on the symmetry. The benefits of the dense *ansatz* are greatest when the Luttinger parameter K is small.

Although more accurate at a fixed bond dimension, the dense *ansatz* requires more computational resources. For high-accuracy calculations ($\delta E/U < 10^{-6}$) the sparse *ansatz* runs faster and uses less memory. At moderate accuracy, however, the dense *ansatz* is more efficient. The threshold value of $\delta E/U$ falls with decreasing K .

Our results are relevant for a wide variety of systems. Any gapless system (including quasi-2D geometries) will invariably have critical Luttinger-liquid-like features when modeled using a MPS. Moreover, our considerations apply to all tensor network approaches [32,33]. Efficient numerical calculations require an awareness of the interplay between spontaneous symmetry breaking and conservation laws.

ACKNOWLEDGMENT

This work was supported by the NSF Grant No. PHY-2110250.

APPENDIX A: IMPLEMENTATION-DEPENDENT METRICS

Here we compare *computer memory usage* and *wall time per iteration* for iDMRG calculations of the ground state of the Bose-Hubbard model, using either a sparse or dense representation of the tensors in the matrix product state. These metrics depend on hardware and implementation details. Here we use the iDMRG algorithm [1,29] using the ITENSOR C++ library [10] compiled with Intel MKL on a single core without multithreading. The code for these calculations involves only minor tweaks to the native iDMRG code on ITENSOR [34]. Despite the implementation dependence, we expect qualitative features to be generic.

The main conclusions are as follows: (1) For small Luttinger parameter K , it is favorable to use the dense MPS *ansatz*, unless one targets an extremely high accuracy. For moderate accuracies, the dense *ansatz* takes less memory and




FIG. 5. Variational energy of the 1D Bose-Hubbard model (E), measured relative to the extrapolated ground-state energy (E_0), as a function of bond dimension χ . Line color corresponds to different values of t/U , resulting in Luttinger parameters $K = 0.42, 0.11$, shown as red or black (cf. inset of Fig. 2). Solid lines: dense *ansatz*, with $N_{\text{tot}} = \chi^2 d$ degrees of freedom, where d is the size of the local Hilbert space ($d = 8$ for $K = 0.11$ and $d = 6$ for $K = 0.42$); dotted-dashed lines: sparse number-conserving *ansatz*, with fewer degrees of freedom N_{tot} . For a fixed χ the dense *ansatz* always yields a smaller error, but the benefit decreases with increasing K . Dotted lines: sparse number-conserving *ansatz*, plotted vs effective bond dimension $\chi_{\text{eff}} \equiv \sqrt{N_{\text{tot}}/d}$, which is a measure of the memory required to store the state. For $K = 0.42$ and a fixed memory footprint, the sparse *ansatz* yields more accurate results. For $K = 0.11$ there is little difference, but the dense *ansatz* is slightly more accurate.

results in a faster calculation. This behavior is analogous to the transverse-field Ising model in the symmetry-broken phase. (2) For large K it is always favorable to use the sparse *ansatz*. This behavior is analogous to the transverse-field Ising model in the paramagnetic phase.

1. Memory usage

Each tensor in our dense MPS *ansatz* requires storing $N_{\text{tot}} = \chi^2 d$ numbers. Here d is the dimension of the local Hilbert space: In the Bose-Hubbard model, $d = n_{\text{max}} + 1$ is set by the maximum number of particles that we allow on a site. For a fixed χ , the sparse representation requires a smaller N_{tot} , as we do not need to store the entries which vanish due to symmetry. To compare the memory usage of the two approaches, we define $\chi_{\text{eff}} = \sqrt{N_{\text{tot}}/d}$. For our Bose-Hubbard calculations we find that for moderate $\chi \lesssim 400$ there is a nearly linear relationship between χ and χ_{eff} .

In Fig. 5 we show the accuracy of the dense iDMRG energy as a function χ (solid curves) and the sparse iDMRG energy as a function of χ_{eff} (dotted curves). For $K = 0.42$ the sparse *ansatz* requires a smaller memory footprint to achieve the same accuracy (the red dotted curve lies below the solid red curve). Conversely, for $K = 0.11$, the sparse *ansatz* requires a larger footprint. These observations are in line with the arguments from Sec. II, which suggest that at




FIG. 6. Variational energy of the 1D Bose-Hubbard model versus wall time per iteration, as bond dimension is varied. Small bond dimensions correspond to short wall times, but low accuracy (see Fig. 5). The number-conserving calculations (dotted-dashed lines) have a steeper slope than those using dense matrices (solid lines), corresponding to a better scaling with bond dimension. Nonetheless, for moderate accuracy calculations $(E - E_0)/U \gtrsim 10^{-6}$, the dense *ansatz* results in faster calculations. The crossover point depends on K , and on implementation details. Smaller K benefits the dense calculation.

small K , where we are proximate to a Bose-Einstein condensate, the dense *ansatz* can more efficiently encode the quantum state. For larger K the advantage goes to the sparse *ansatz*.

Figure 5 also plots the sparse data versus χ as dotted-dashed lines which always lie above the solid lines: For a fixed χ , the dense *ansatz* has more degrees of freedom, and hence yields a lower variational energy. As expected, the advantage is greatest for small K .

2. Wall time

In Fig. 6 we plot the same variational energies shown in Fig. 5, but now versus the wall time per iteration. As previously noted, these times are highly dependent on implementation, and the actual number is not particularly relevant. Nonetheless, we can use these timings to compare the performance of the two *ansätze*.

For the dense *ansatz* with $K = 0.11$, one achieves an accuracy of $(E - E_0)/U = 10^{-5}$ by taking $\chi = 19$. For these parameters, each iteration of the iDMRG algorithm takes a time of 0.1 s. By contrast, achieving the same accuracy with the sparse *ansatz* requires a larger $\chi = 87$, and each iteration takes significantly longer, 3.2 s. The sparse algorithm, however, scales much better with bond dimension, and if one has a target accuracy of $(E - E_0)/U = 10^{-7}$, the two calculations take the same time. If higher accuracy is required, the sparse *ansatz* is faster.

These features are illustrated in Fig. 6 by the steeper slope of the sparse data. As one moves to larger K , the relative performance of the sparse *ansatz* improves. In particular, at

$K = 0.42$, it is more time efficient to use the sparse *ansatz* if one requires an accuracy smaller than $(E - E_0)/U = 10^{-6}$. Increasing K farther continues to move this crossover point to lower accuracy.

APPENDIX B: iDMRG DETAILS

In our calculations we start with a product state, then implement the iDMRG algorithm with two-site updates to find MPS approximations of the ground state in the thermodynamic limit. Simulations are carried out using the ITENSOR library [10]. Here we describe several technical details.

1. Truncation error

For calculations of the transverse-field Ising model in Sec. III A, we increase the bond dimension as necessary until the properties of the state have all converged. We find that it is sufficient to reduce the truncation error $e_{\text{trc}} \leq 10^{-12}$ to achieve convergence. One can interpret the data in Fig. 1 as numerically exact results.

As for the Bose-Hubbard (Sec. III B) and Fermi-Hubbard (Sec. III C) simulations, “convergence” is no longer a meaningful criterion. The ground states are gapless critical states and the long-distance properties of the correlation functions cannot be modeled by matrix product states with fixed bond dimension. As has been argued elsewhere [22,23], however, features of the asymptotic ground state can be inferred by studying the behavior of variational wave functions as a function of the bond dimension. This is known as finite-entanglement scaling. For the data shown in Figs. 2 and 3, we increase the bond dimension from $\chi = 40$ to 120 in steps of $\Delta\chi = 10$. We find that the truncation error scales as a power law of the bond dimension $e_{\text{trc}} \propto \chi^{-2\kappa}$, consistent with the Luttinger-liquid scaling observed in other observables [23].

2. Fixing the chemical potential

Throughout this paper, we calculate properties at fixed density. This constraint is simple to incorporate into the sparse number-conserving MPS *ansatz*. For the dense *ansatz*, one instead has to specify a suitable chemical potential μ to fix the density at the desired value.

To achieve our target density $\bar{n}^*$, we vary the chemical potential in the early iterations of the dense iDMRG algorithm. The update procedure involves approximating the inverse compressibility based on measurements made in subsequent iterations:

$$\frac{\partial \mu_i}{\partial \bar{n}_i} \approx \frac{\mu_i - \mu_{i-1}}{\bar{n}_i - \bar{n}_{i-1}}. \quad (\text{B1})$$

We then define the chemical potential for iteration $i + 1$ based on the compressibility computed in iteration i :

$$\mu_{i+1} = \mu_i + \alpha_i(\bar{n}^* - \bar{n}_i) \frac{\partial \mu_i}{\partial \bar{n}_i}. \quad (\text{B2})$$

Here the convergence factor α controls how large of an update we allow from iteration to iteration. Smaller α results in a more stable algorithm, but slower convergence. In all our calculations we take $\alpha_i = 0.1$.

APPENDIX C: SCALING ANALYSIS

Here we describe how we use a scaling analysis of the single-particle density matrix to extract the Luttinger parameter for the Bose-Hubbard model. We use this analysis in Secs. III B and IV.

As argued in the main text, the correlation length of the MPS is expected to scale as χ^κ where $\kappa = 6/(1 + \sqrt{12})$ [23]. For distances that are smaller than this correlation length, we expect $\langle a_i^\dagger a_j \rangle \propto |i - j|^{-\kappa/2}$. Given a guess for the optimal

Luttinger parameter K_0 , we rescale

$$|i - j| \rightarrow |i - j|\chi^{-\kappa},$$

$$\langle a_i^\dagger a_j \rangle \rightarrow \langle a_i^\dagger a_j \rangle \chi^{\kappa K_0/2}.$$

Following a procedure similar to Ref. [22], we then define an objective function which measures the deviation between the scaled density matrices with different values of χ . These should all collapse when $K_0 = K$. We adjust K_0 to minimize our objective function.

For the case in Sec. III B, where $\bar{n} = 1$ and $U/t = 3$, we find $K = 0.423(2)$. Note that the value of the Luttinger parameter at the Mott lobe tip is 0.5, so this result is consistent with being on the superfluid side of the Mott-superfluid transition.

This same bond dimension scaling procedure is used to generate the data in Fig. 4.

-
- [1] U. Schollwöck, *Ann. Phys.* **326**, 96 (2011), January 2011 Special Issue.
 - [2] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, *Ann. Phys.* **411**, 167998 (2019).
 - [3] E. Noether, Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse **1918**, 235 (1918).
 - [4] E. Noether, *Transp. Theory Stat. Phys.* **1**, 186 (1971).
 - [5] A. Beekman, L. Rademaker, and J. van Wezel, *SciPost Phys. Lect. Notes*, doi:10.21468/scipostphyslectnotes.11 (2019).
 - [6] N. D. Mermin and H. Wagner, *Phys. Rev. Lett.* **17**, 1133 (1966).
 - [7] P. C. Hohenberg, *Phys. Rev.* **158**, 383 (1967).
 - [8] T. Momoi, *J. Stat. Phys.* **85**, 193 (1996).
 - [9] T. Giamarchi, *Quantum Physics in One Dimension* (Oxford University Press, Oxford, 2003).
 - [10] M. Fishman, S. R. White, and E. M. Stoudenmire, *arXiv:2007.14822*.
 - [11] L. Pitaevskii and S. Stringari, *Bose-Einstein Condensation and Superfluidity* (Oxford University Press, Oxford, 2016).
 - [12] I. P. McCulloch, *J. Stat. Mech.: Theory Exp.* **2007**, P10014 (2007).
 - [13] P. Pfeuty, *Ann. Phys.* **57**, 79 (1970).
 - [14] P. Pfeuty and R. J. Elliott, *J. Phys. C: Solid State Phys.* **4**, 2370 (1971).
 - [15] P. Coleman, *Introduction to Many-Body Physics* (Cambridge University Press, Cambridge, 2015).
 - [16] U. Schollwöck, *Rev. Mod. Phys.* **77**, 259 (2005).
 - [17] D. M. Greenberger, M. A. Horne, and A. Zeilinger, Going beyond bell's theorem, in *Bell's Theorem, Quantum Theory and Conceptions of the Universe*, edited by M. Kafatos (Springer, Dordrecht, 1989), pp. 69–72.
 - [18] H.-C. Jiang, Z. Wang, and L. Balents, *Nat. Phys.* **8**, 902 (2012).
 - [19] H.-C. Jiang and L. Balents, *arXiv:1309.7438* (2013).
 - [20] F. D. M. Haldane, *Phys. Rev. Lett.* **47**, 1840 (1981).
 - [21] M. A. Cazalilla, R. Citro, T. Giamarchi, E. Orignac, and M. Rigol, *Rev. Mod. Phys.* **83**, 1405 (2011).
 - [22] T. G. Kiely and E. J. Mueller, *Phys. Rev. B* **105**, 134502 (2022).
 - [23] F. Pollmann, S. Mukerjee, A. M. Turner, and J. E. Moore, *Phys. Rev. Lett.* **102**, 255701 (2009).
 - [24] H. Li and F. D. M. Haldane, *Phys. Rev. Lett.* **101**, 010504 (2008).
 - [25] R. Thomale, D. P. Arovas, and B. A. Bernevig, *Phys. Rev. Lett.* **105**, 116805 (2010).
 - [26] F. Pollmann, A. M. Turner, E. Berg, and M. Oshikawa, *Phys. Rev. B* **81**, 064439 (2010).
 - [27] E. H. Lieb and F. Y. Wu, *Phys. Rev. Lett.* **20**, 1445 (1968).
 - [28] H. J. Schulz, *Phys. Rev. Lett.* **64**, 2831 (1990).
 - [29] I. P. McCulloch, *arXiv:0804.2509*.
 - [30] D. M. Gangardt and G. V. Shlyapnikov, *Phys. Rev. Lett.* **90**, 010401 (2003).
 - [31] H. F. Song, S. Rachel, and K. Le Hur, *Phys. Rev. B* **82**, 012405 (2010).
 - [32] B. Bauer, P. Corboz, R. Orús, and M. Troyer, *Phys. Rev. B* **83**, 125106 (2011).
 - [33] S. Singh, R. N. C. Pfeifer, and G. Vidal, *Phys. Rev. B* **83**, 115125 (2011).
 - [34] E. M. Stoudenmire, <https://github.com/ITensor/iDMRG>.