Finite projected entangled pair states for the Hubbard model

M. Scheb and R. M. Noack

Fachbereich Physik, Philipps-Universität Marburg, 35032 Marburg, Germany

We adapt and optimize the projected-pair-entangled-state (PEPS) algorithm on finite lattices (fPEPS) for two-dimensional Hubbard models and apply the algorithm to the Hubbard model with nearest-neighbor hopping on a square lattice. In particular, we formulate the PEPS algorithm using projected entangled pair operators, incorporate SU(2) symmetry in all tensor indices, and optimize the PEPS using both iterative-diagonalization-based local bond optimization and gradient-based optimization of the PEPS. We discuss the performance and convergence of the algorithm for the Hubbard model on lattice sizes of up to 8×8 for PEPS states with U(1) symmetric bond dimensions of up to $D = 8$ and SU(2) symmetric bond dimensions of up to $D = 6$. Finally, we comment on the relative and overall efficiency of schemes for optimizing fPEPS.

I. INTRODUCTION

One of the most important models used to study strongly correlated electron systems is the Hubbard model [1]. Originally, it was conceived to model narrow energy bands within many-body quantum systems in order to study the behavior of observables at low temperatures and, in particular, to study the Mott transition [2,3]. With the discovery of high-temperature superconductors in the cuprates in 1986 [4], the focus of interest shifted to using the two-dimensional Hubbard model to model the electronic transport within copper-dioxide planes common to all the high-T_c cuprates [5–8]. A tendency towards superconducting pairing with d-wave symmetry found in numerical and in perturbation-theory calculations [9] reinforced the picture of the two-dimensional Hubbard model as a minimal model for high-T_c superconductivity in the cuprates. Experimental studies by Tranquada and coworkers [10–12] further expanded knowledge of phenomena found in cuprates in that stripe structures, consisting of oscillating spins and charge densities, were found to be present in the underdoped region in a variety of high-T_c cuprates. These stripe phases generally compete with the superconducting pairing, so that superconductivity is suppressed, especially at particular commensurate band fillings. Subsequent studies of the t-J-model [13] and the underdoped Hubbard model [14] were able to qualitatively reproduce aspects of this characteristic behavior.

Despite these and many other insights gained into the behavior of the two-dimensional Hubbard model within and outside of the context of high-temperature superconductors, many aspects of the low-temperature behavior remain contentious, such as the competition of different pairing orders [14], the proper description of topological phases, and the overall behavior in the thermodynamic limit in two dimensions [15,16].

Many analytical and numerical methods exist for studying strongly correlated quantum systems in general and the Hubbard model in particular. For the ground-state calculation in one dimension, the most efficient numerical method is the density matrix renormalization group (DMRG) [17,18],

in which the wave function is efficiently approximated as a matrix product state (MPS) [19]. Coupled chains as well as more general two-dimensional systems of limited width have been studied extensively using the DMRG [14,20–34] by mapping the two-dimensional lattice onto the intrinsically one-dimensional MPS. However, due to the linear increase of entropy and thus exponential increase of computational effort with the width of the lattice in two dimensions, the thermodynamic limit will probably remain inaccessible [35,36]. In order to circumvent this exponential computational hurdle, Verstraete *et al.* introduced projected entangled pair states (PEPS) [37], which are the natural generalization of MPSs to two dimensions. It has been shown that PEPSs are the ground states of local Hamiltonians [38–40] and are thus promising candidates for studying two-dimensional quantum systems. Murg *et al.* conducted one of the first numerical simulations and provided benchmark results for hard-core bosons [41] and frustrated spin systems [42]. Further PEPS-based simulations were conducted by Lubasch and coworkers [43,44], who focused on the Heisenberg model.

Development and application of methods to treat fermions using PEPS has been led by Corboz and coworkers, who have concentrated primarily on the infinite PEPS (iPEPS) algorithm originally introduced by Jordan *et al.* [45]. In this method, translational invariance is emulated by replicating a unit cell infinitely in all spatial dimensions. To distinguish the original version of PEPS with open boundary conditions from iPEPS, we designate it finite PEPS (fPEPS) here. In a series of papers [46–49], Corboz *et al.* used both fPEPS and iPEPS to examine the t-J-model, a simplified version of the Hubbard model. They found the stripe structure, in which both the charge and the spin density of electrons oscillate over multiple sites and compete with d-wave pairing. In recent years, the most prominent variation of PEPS algorithms has been iPEPS optimized via imaginary time evolution. In addition, gradient optimizations of the entire PEPS have proven to be successful [50–52].

In this paper, we revisit the original finite PEPS algorithm, which treats a finite lattice with open boundary conditions.

Working on a specific finite lattice has the decided advantage that energies and observables can be compared directly to those obtained with other finite-lattice numerical methods, in particular, with exact diagonalization and with the DMRG, both of which are strictly variational. We note that PEPS-based algorithms are not strictly variational in general because the contraction of a two-dimensional tensor network must be carried out approximately in order to be efficient [37,40]. In addition, optimization of PEPS states has typically been carried out using imaginary-time evolution within the Trotter approximation, which introduces an additional systematic Trotter error, further complicating a variational comparison. Therefore, our comparison will give us a stringent test bed for evaluating the accuracy and convergence behavior of the fPEPS as well as determining to what extent it obeys the variational principle.

In developing our variant of the fPEPS algorithm, we have made an effort to incorporate as large a scope of modern ideas and methods as possible. In particular, the fact that MPS-based algorithms have profited significantly from the formulation of matrix-product operators (MPOs) has motivated us to develop a general scheme for formulating fPEPS algorithms in terms of the generalization of MPOs, projected entangled pair operators (PEPOs) [53]. Since the Hubbard model at general band filling and zero magnetic field has SU(2) symmetry in the spin sector and U(1) symmetry in the charge sector, we incorporate these symmetries explicitly into our construction of PEPSs and PEPOs. In addition, we build bookkeeping for the fermionic sign into our PEPOs, which is a different formulation than the original one of Refs. [46,54] in terms of fermionic swap gates. For MPS algorithms, variational iterative diagonalization of a local effective Hamiltonian is a powerful tool for optimizing an MPS and forms the basis of the DMRG algorithm. We adapt a scheme for carrying out a local-bond-based variational optimization of a PEPS [55] to fPEPS, formulating variants that carry out single site as well as bond optimization. Note that a local-optimization-based treatment of the Hubbard model within fPEPS using the a scheme based on Ref. [56] was recently described in Ref. [57]. In order to reduce the size of the effect Hilbert spaces that must be treated within local variational optimization, we generalize the recently formulated controlled bond expansion (CBE) of Gleis *et al.* [58] to reduce the numerical cost of local bond optimization in fPEPS. We also apply an alternative optimization scheme for PEPS based on a gradient of the energy functional that was first applied by Vanderstraeten *et al.* [50] to PEPS algorithms to fPEPS. In addition, we have reexamined and reworked the scheme for approximately contracting an fPEPS network first described in Ref. [37], concentrating on improving stability and efficiency; here we have found a scheme to reduce the numerical effort needed to optimize the contracted environment using an adaptation of the CBE mentioned above [58].

Clearly, our method is tailored for short-ranged two-dimensional strongly interacting quantum lattice models, especially ones with fermionic degrees of freedom and SU(2) spin symmetry. Thus, our method is intended to treat the two-dimensional Hubbard model, its extensions, as well as other related models (e.g., Heisenberg models). As a benchmark system, we take the usual two-dimensional Hubbard model

with nearest-neighbor hopping on a finite square lattice with open boundary conditions, i.e., treat the Hamiltonian

$$\mathcal{H} = -t \sum_{\langle i,j \rangle, \sigma} (c^\dagger_{i,\sigma} c_{j,\sigma} + \text{H.c.}) + U \sum_i n_{i,\uparrow} n_{i,\downarrow}, \quad (1)$$

where

$$n_{i,\sigma} = c^\dagger_{i,\sigma} c_{i,\sigma}$$

is the local particle-density operator. The notation $\langle i, j \rangle$ indicates nearest-neighbor sites on an $L \times L$ lattice with open boundary conditions. The total number of sites is then $V = L^2$. We work in the canonical ensemble, so that the numbers of spin-up electrons $N_\uparrow$ and spin-down electrons $N_\downarrow$ or, equivalently, the total particle number $N \equiv N_\uparrow + N_\downarrow$ and the z component of the total spin $S_z \equiv (N_\uparrow - N_\downarrow)/2$, are conserved. Conservation of total particle number corresponds also to the conservation of the deviation from half filling $C_z \equiv (N_\uparrow + N_\downarrow - V)/2$, which is the charge-sector analog of S_z. Here we will be interested in non-magnetized ground states, so we will take $S_z = 0$ in this paper. Ground states in the $S_z = 0$ sector generically have total spin $S = 0$, so that we will concentrate on this case for calculations that explicitly take the SU(2) spin symmetry into account. We additionally specify the overall band filling using the average particle number $\langle n \rangle = N/V$ and take $t = 1$ to be the scale of the Hamiltonian so that there are two independent physical parameters, the on-site interaction U and the average particle number $\langle n \rangle$.

This paper is organized as follows: Section II A gives a concise definition of PEPSs and their entanglement structure. In Sec. II B, we construct the corresponding projected entangled pair operator (PEPO), which is the Hamiltonian adapted to the topology of the PEPS. Section III introduces a generic scheme to incorporate SU(2)-symmetries into the PEPS-PEPO framework. In Sec. IV A, we explain in detail how to calculate expectation values approximately within the PEPS-PEPO scheme. Section IV B presents two different optimization procedures; the first is comprised of local updates implemented using a sequence of iterative diagonalizations of local effective Hamiltonians, similar to that used in the DMRG. The second procedure is based on direct gradient optimization of the PEPS tensors. We find that a combination of the two procedures provides the best results. In Sec. V, we present benchmark results for the two-dimensional Hubbard model, discussing accuracy and the convergence and its dependence on overall bond dimension and on environment dimension for lattices sizes ranging from 3×3 to 8×8. Finally, we discuss the convergence issues and the state of development of the fPEPS-PEPO scheme in Sec. VI.

II. PROJECTED ENTANGLED PAIR STATE SCHEME

In this section, we will define the states and operators used in PEPS-based methods, concentrating on the finite-lattice PEPS algorithm, as first proposed in Ref. [37]. We use a tensor-network formulation to define operators and apply them to a PEPS, i.e., we work in terms of PEPOs, analogous to matrix product operators (MPOs). The extension of an MPO to a more general tensor network can be termed a tensor network operator (TNO), of which a PEPO is a special case.
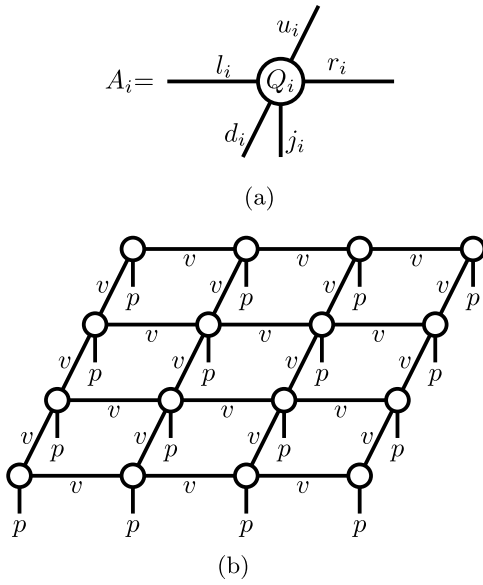
FIG. 1. Projected entangled pair state. (a) PEPS bulk tensor and (b) PEPS on a 4×4 lattice.

We will first give a general definition of an fPEPS, followed by a definition of a PEPO and its construction.

A. Finite projected entangled pair state

Consider a generic many-body wave function on N sites

$$|\psi\rangle = \sum_{j_1,\dots,j_N} \psi_{j_1,\dots,j_N} |j_1,\dots,j_N\rangle. \quad (2)$$

Here $|j_1,\dots,j_N\rangle = |j_1\rangle \otimes \cdots \otimes |j_N\rangle$ are many-body basis states with $j_i = 1,\dots,d$, where d is the dimension of a local Hilbert space. A convenient way to store this wave function on a two-dimensional lattice is as a PEPS as introduced by Verstraete and Cirac in 2004 [37]. In a PEPS, each bond between adjacent sites is associated with a maximally entangled state

$$|\phi\rangle = \sum_{k=1}^{D} |k,k\rangle \quad (3)$$

in a virtual Hilbert space, where D is the maximum bond dimension. Conversely, this means that a bulk site i is connected to four virtual bonds u_i, d_i, l_i, and r_i. If we define a projector Q_i that maps virtual bonds at site i to the physical bond j_i, we can define the rank-5 PEPS tensor

$$A_i = A_{u_i,d_i,l_i,r_i}^{j_i} = \langle j_i|Q_i|u_i,d_i,l_i,r_i\rangle, \quad (4)$$

which is depicted graphically [59] in Fig. 1(a). Tensors at the edges or the corners of the lattice are rank-4 and rank-3, accordingly. There is no straightforward way of writing the entire PEPS as an explicit equation due to the two-dimensional arrangement of tensors, which is why we define it simply as

$$|\psi\rangle = \sum_{j_1,\dots,j_N} \mathcal{F}(A_1 A_2 \dots A_N)|j_1,\dots,j_N\rangle,$$
$$\mathcal{F}(A_i A_2 \dots A_N) = \psi_{j_1,\dots,j_N} \quad (5)$$

for a lattice of N sites, where $\mathcal{F}$ is a function that contracts common indices in PEPSs. Figure 1(b) illustrates this arrangement on a 4×4 lattice with virtual indices v and physical indices p. Note that the construction of A_i via Q_i is mainly academic; in practical simulations, a PEPS will be initialized as a product state with virtual indices of dimension one or via random tensors. As the optimization proceeds, the dimensions then may change and lead to states, which are, in general, not maximally entangled.

PEPSs are promising candidates for describing ground states in two-dimensional quantum systems [39,40,44,55, 60,61]. Their introduction, however, leads to two complications: First, the exact calculation of expectation values $\langle \psi|O|\psi\rangle$ scales exponentially with system size, which is why it has to be carried out approximately, and, second, due to the high rank of bulk tensors, PEPS-based algorithms scale with a high power of the number of virtual states D. We will discuss these complications as well as ways to deal with them in more detail in the following.

B. Projected entangled pair operators

Consider an electronic many-body system on a two-dimensional square lattice that is described by the Hubbard Hamiltonian (1), which we rewrite slightly as
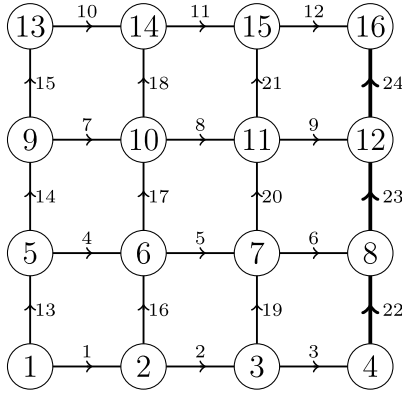
$$\mathcal{H} = -t \sum_{\langle i,j\rangle,\sigma} (c_{i,\sigma}^\dagger c_{j,\sigma} + \text{H.c.}) + U \sum_i c_{i,\uparrow}^\dagger c_{i,\downarrow}^\dagger c_{i,\downarrow} c_{i,\uparrow}. \quad (6)$$

The fermionic creation and annihilation operators can be written in terms of their action on the four basis states $\{|0\rangle, |\uparrow\rangle, |\downarrow\rangle, |\uparrow\downarrow\rangle\}$ spanning the local Hilbert space,

$$c_\uparrow^\dagger = |\uparrow\rangle\langle 0| + |\uparrow\downarrow\rangle\langle\downarrow|,$$
$$c_\uparrow = |0\rangle\langle\uparrow| + |\downarrow\rangle\langle\uparrow\downarrow|,$$
$$c_\downarrow^\dagger = |\downarrow\rangle\langle 0| - |\uparrow\downarrow\rangle\langle\uparrow|,$$
$$c_\downarrow = |0\rangle\langle\downarrow| - |\uparrow\rangle\langle\uparrow\downarrow|.$$

We wish to store the expression corresponding to Eq. (6) as a PEPO, which has the same topology as the PEPS to which it is supposed to be applied. We start by expanding the full Hamiltonian for a lattice of a given width and height. For example, for a 4×4 lattice with open boundary conditions, the full operator reads, schematically,

$$\mathcal{H} = (Uc_\uparrow^\dagger c_\downarrow^\dagger c_\downarrow c_\uparrow)_1 \otimes \mathbb{1}_2 \otimes \mathbb{1}_3 \otimes \cdots \otimes \mathbb{1}_{16}$$
$$+ \mathbb{1}_1 \otimes (Uc_\uparrow^\dagger c_\downarrow^\dagger c_\downarrow c_\uparrow)_2 \otimes \mathbb{1}_3 \otimes \cdots \otimes \mathbb{1}_{16}$$
$$\cdots$$
$$+ \mathbb{1}_1 \otimes \mathbb{1}_2 \otimes \cdots \otimes \mathbb{1}_{15} \otimes (Uc_\uparrow^\dagger c_\downarrow^\dagger c_\downarrow c_\uparrow)_{16}$$
$$+ (-tc_\uparrow^\dagger \mathcal{P})_1 \otimes (c_\uparrow)_2 \otimes \mathbb{1}_3 \otimes \cdots \otimes \mathbb{1}_{16}$$
$$+ \mathbb{1}_1 \otimes (-tc_\uparrow^\dagger \mathcal{P})_2 \otimes (c_\uparrow)_3 \otimes \mathbb{1}_4 \otimes \cdots \otimes \mathbb{1}_{16}$$
$$\cdots$$
$$+ (-tc_\uparrow^\dagger \mathcal{P})_1 \otimes \mathcal{P}_2 \otimes \mathcal{P}_3 \otimes \mathcal{P}_4 \otimes (c_\uparrow)_5$$
$$\otimes \mathbb{1}_6 \otimes \cdots \otimes \mathbb{1}_{16}$$
$$\cdots$$

FIG. 2. Sites and signaling channels in between for a 4×4 lattice.

where the parity operator

$$\mathcal{P} = |0\rangle\langle 0| - |\uparrow\rangle\langle\uparrow| - |\downarrow\rangle\langle\downarrow| + |\uparrow\downarrow\rangle\langle\uparrow\downarrow|$$

encodes the effect of the fermionic sign. The subscripts denote the physical sites on which the operators act, whereas the enumeration is chosen according to the scheme depicted in Fig. 2.

Inspired by the construction of matrix product operators via finite-state machines (FSMs), which was introduced by Crosswhite *et al.* [62,63] and refined by Paeckel *et al.* [64], we now conceptualize $\mathcal{H}$ as a regular expression. In particular, for each unique sequence of operators on one site, we substitute one symbol ($U c_\uparrow^\dagger c_\downarrow^\dagger c_\downarrow c_\uparrow \to$ U, $-t c_\uparrow^\dagger \mathcal{P} \to$ A, $c_\uparrow \to$ B, $-t c_\downarrow^\dagger \mathcal{P} \to$ C, $c_\downarrow \to$ D, $t c_\uparrow \mathcal{P} \to$ E, $c_\uparrow^\dagger \to$ F, $t c_\downarrow \mathcal{P} \to$ G, $c_\downarrow^\dagger \to$ H, $\mathbb{1} \to$ I, $\mathcal{P} \to$ P). Nearest-neighbor hopping terms are thus described by the four pairs AB, CD, EF, and GH. Tensor products connecting different local Hilbert spaces are then interpreted as concatenations, whereas the sums become unions. In this way, all symbols form an alphabet, each summand of the many-body Hamiltonian becomes a word, and the set of all words forms a regular language,

$$H = \{\text{UIIIIIIIIIIIIIII}, \text{IUIIIIIIIIIIIIII}, \dots, \text{IIIIIIIIIIIIIIIU},$$
$$\text{ABIIIIIIIIIIIIII}, \text{IABIIIIIIIIIIIII}, \dots,$$
$$\text{APPPBIIIIIIIIIII}, \dots\}.$$

Here we drop the subscripts—they are determined by the position of a symbol in the sequence.

The PEPO we want to form will be a set of interconnected FSMs, which generate all of the words of the language in a two-dimensional fashion. In order to do this, we connect adjacent sites through directed signaling channels, as depicted in Fig. 2. The flow of information is defined as going upwards and to the right. Information at the upper boundary is dropped, while information at the right boundary is redirected to the top, which makes the rightmost vertical channels of the lattice the trunk. The upper-right corner (site 16 in Fig. 2) is thus the sink, i.e., the location at which all the information passed by the FSMs is gathered. Here it is determined whether a string of symbols is actually a part of the language and, hence, is a meaningful term of the Hamiltonian. For example, channel 5 reports what symbols have occurred on sites 5 and 6, and channel 21 what symbols have occurred on sites 3, 7, and 11.

TABLE I. Finite-state machine for site 6 implemented as an associative tensor. For the incoming states $s_{4,i}$ from channel 4 and $s_{16,i}$ from channel 16, the symbol σ_i is inserted into the current word, and $s_{5,i}$ and $s_{17,i}$ are emitted via channels 5 and 17.

C 4	C 16	C 5	C 17	S 6
$s_{4,1}$	$s_{16,1}$	$s_{5,1}$	$s_{17,1}$	σ_1
$s_{4,2}$	$s_{16,2}$	$s_{5,2}$	$s_{17,2}$	σ_2
$s_{4,3}$	$s_{16,3}$	$s_{5,3}$	$s_{17,3}$	σ_3
$\dots$				

Channel 23, as a trunk channel, carries information from the entire block below, i.e., from all sites from 1 to 8. In this way, each FSM can be seen as an associative tensor or table in which a set of incoming and outgoing states forms a key, the associated symbol represents their value, and the pair of both is an element. The rank is equal to the coordination number, which is two at the corners, three at the edges, and four otherwise. As an example, the FSM on site 6 is schematically described in Table I.

We define three distinct states, s_i, s_f, and s_P, which designate the initial state, the final state, and the parity state, respectively. The symbol s_i indicates that, up to this point, only identities (designated as "I") have appeared, and the first occurrence of a nontrivial symbol is pending. The symbol s_f designates the complement, namely that a valid combination of nontrivial symbols has already appeared and, after the current point, only identity symbols are allowed to be attached to the word. Finally, the symbol s_P takes the parity operators (designated as "P") from the Jordan-Wigner transformation into account and connects vertical hopping terms in a way that will be elucidated via an example later on. These states, together with the two trivial symbols "I" and "P", are used to initialize the PEPO as depicted in Tables II and III. Note how there are no elements with two incoming final states s_f, which prevents different words from mixing with each other.

After all of these preparations, the Hamiltonian can now be converted into a PEPO by recasting every word into a set of nontrivial pairs, with the first element being a nontrivial symbol and the second the index of the site it acts on. Each pair is then inserted into the PEPO by adding an element to the tensor of the respective site. In this way, the entire PEPO can be systematically constructed word for word.

For example, the word "IIIIIUIIIIIIIIII" translates as $\{(U, 6)\}$ and is incorporated into the PEPO by adding the element $\{(s_i, s_i, s_f, s_f), U\}$ to the tensor at site 6 in accordance

TABLE II. Initial elements of a bulk tensor/FSM. Here C I_l and C I_b are the two incoming channels from left and bottom, and C O_r and C O_t are the outgoing channels to the right and top, respectively.

C I_l	C I_b	C O_r	C O_t	S
s_i	s_i	s_i	s_i	I
s_i	s_f	s_i	s_f	I
s_f	s_i	s_f	s_i	I
s_i	s_P	s_i	s_P	I
s_f	s_P	s_f	s_P	I
s_P	s_i	s_P	s_P	P

TABLE III. Initial elements of a trunk tensor/FSM. Here C I_l and C I_b are the two incoming channels from left and bottom, and O_t is the outgoing channel to the top.

C I_l	C I_b	C O_t	S
s_i	s_i	s_i	I
s_f	s_i	s_f	I
s_i	s_f	s_f	I
s_f	s_P	s_f	I
s_P	s_i	s_P	P

with Table I. The remaining 15 identity symbols not explicitly contained in the set are taken care of by the previous initialization.

The next type of words are those which represent horizontal hopping terms, such as "IIIIIABIIIIIIIII", as is depicted in Fig. 3. The set of pairs for this case reads $\{(A, 6), (B, 7)\}$ and the elements, which need to be attached to the tensors of site 6 and 7 are $\{(s_i, s_i, \{(A, 6)\}, s_f), A\}$ and $\{(\{(A, 6)\}, s_i, s_f, s_f), B\}$, respectively. This means that at site 6, a word is initialized by inserting the symbol "A". A final state is emitted upwards, meaning that no other nontrivial symbol is expected above, while the intermediate state $\{(A, 6)\}$ is sent to the right. At site 7, the latter state is received from the left, the second symbol B is attached to the word, and a final state is emitted both upwards and to the right.

Finally, we consider terms that span multiple rows, taking "IAPPPBIIIIIIIII", which describes vertical hopping between sites 2 and 6, as a concrete example; this case is depicted in Fig. 4. Since the parity operators have been already taken into account through proper initialization, the term is effectively local and can be represented by inserting just two pairs $\{(A, 2), (B, 6)\}$ and the according tensor elements $\{(s_i, s_P, \{(A, 2)\}), A\}$ and $\{(s_P, \{(A, 2)\}, s_f, s_f), B\}$ at sites 2 and 6, respectively.

When a Hamiltonian contains nonlocal terms (a case that we will not further cover explicitly in this paper), elements may have been generated by a previous update of another (nonlocal) term. In this case, a single element can be used as a component of two or more words. Finally, it is crucial to note that intermediate states must be defined as unordered sets, meaning that $\{(A, i), (B, j)\} = \{(B, j), (A, i)\}$.
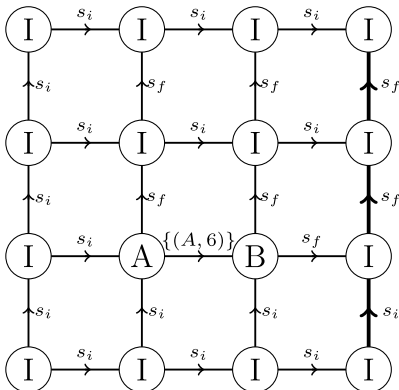

FIG. 3. Graphical depiction of the word "IIIIIABIIIIIIIII" embedded in the PEPO.
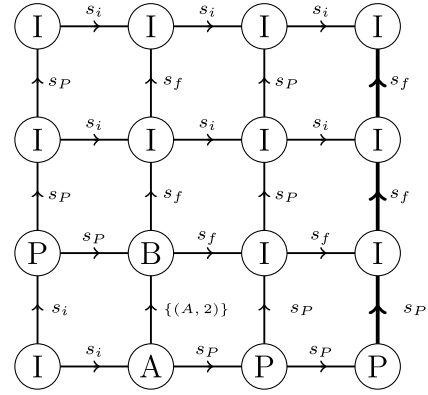

FIG. 4. Graphical depiction of the word "IAPPPBIIIIIIIII" embedded in the PEPO.

The three examples considered above should cover all possible variants of words needed for local Hamiltonians and describe how to insert them into the appropriate tensors. After the PEPO has been fully assigned, all of the intermediate states within keys, which are symbolic, can be encoded as unique numbers so that the PEPO can be stored efficiently. (For the case of the Hubbard model, the number of states required to represent the Hamiltonian is exactly seven, consisting of s_i, s_f, s_p, and the four nontrivial operators.) This numbering can then be translated into the corresponding quantum numbers, which for the Hubbard model treated here are the spin S, its z component S_z, and the deviation of the particle number from half filling C_z. Once these steps have been completed, the symbols within the tensors are replaced by the full quantum-mechanical operators that they represent. To take full advantage of the SU(2)-spin symmetry, the PEPO tensors must be compactified using Eq. (16). The PEPO is now ready to be used to implement the tensor-network algorithm of one's choice, such as variational optimization or imaginary- or real-time evolution.

Finally, we comment on the role of the fermionic sign. Our approach starts by applying the Jordan-Wigner transformation and effectively pushing the fermionic sign into the Hamiltonian through a sequence of parity operators. This allows us to treat all subsequent tensor operations bosonically. Alternatively, one could omit using parity operators altogether and incorporate the sign into the PEPS through fermionic swap gates [46,54], which requires bookkeeping that tracks parity in addition to charge and spin. While this approach seems to be well suited when working with Trotter gates, we have found swap gates to be too cumbersome to use in conjunction with PEPOs. As is evident from our construction, the additional state, s_P, ensures that the Hamiltonian remains local and that the bond dimension of the PEPO does not increase with system size. Note that the PEPO of any local Hamiltonian, bosonic or fermionic, requires a special state that connects vertical hopping terms.

III. SU(2)-SYMMETRIC TENSORS

In order to optimize PEPSs and calculate their expectation values, we need to work efficiently with the tensors of which they and their corresponding PEPOs are composed. Especially
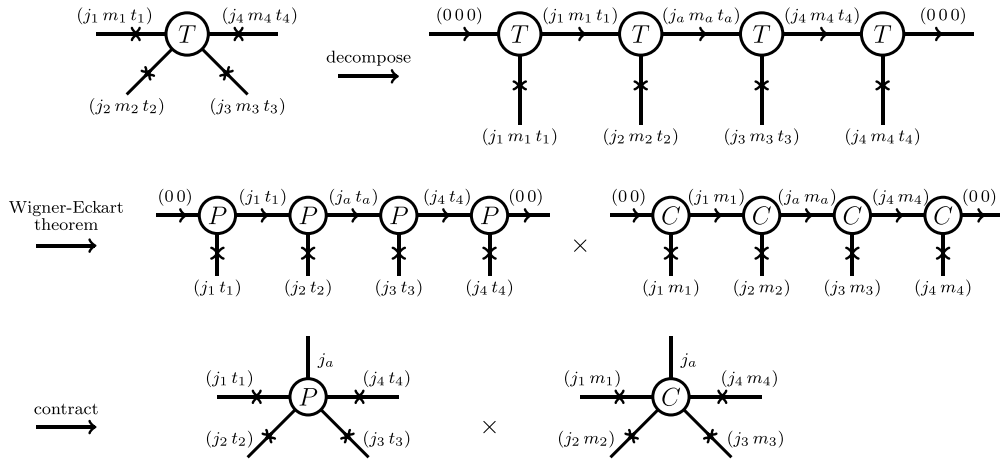
FIG. 5. Compactification of an SU(2)-invariant rank-4 tensor.

useful for this purpose is the exploitation of continuous symmetries, which are present in most local Hamiltonians. The Hubbard model, in particular, exhibits a U(1) symmetry for the charge degrees of freedom and a SU(2) symmetry for the spin degrees of freedom in the absence of an external magnetic field. The implementation of U(1) symmetries in tensor networks has already been covered thoroughly, for instance in Refs. [65,66], and will not be discussed in this paper. Instead, we give a generic and concise documentation of how to incorporate SU(2) symmetries into a networks of PEPS tensors. Some of our concepts are similar to those presented in previous work on SU(2) [67–69], but several techniques differ in ways that may impact both implementation and performance.

In the following, we first define SU(2)-symmetrized tensors in Sec. III A, then define additional manipulations of the tensors and their indexes that we will need to evaluate and optimize PEPS/PEPO tensor networks: permutation (Sec. III B), contraction (Sec. III C), charge fusion (Sec. III D), and index reversal (Sec. III E).

A. Definition

Consider the tensor operator T_q^k, with total angular momentum k and its z component $q \in \{-k, -k+1, \ldots, k-1, k\}$. Consider in particular its transformation properties under the rotation

$$R = R(\boldsymbol{\theta}) = \exp\left(-i\,\boldsymbol{\theta} \cdot \boldsymbol{j}\right), \qquad (7)$$

where $\boldsymbol{\theta}$ is a vectorial angle, and the components of $\boldsymbol{j}$ satisfy the angular-momentum algebra

$$[j_k, j_l] = i\,\epsilon_{klm}\,j_m. \qquad (8)$$

If

$$R\,T_q^k\,R^{-1} = \sum_{q'} T_{q'}^k\,R_{q',q}^k, \qquad (9)$$

with

$$R_{q',q}^k = \langle k\,q' | R | k\,q \rangle, \qquad$$

we can apply the Wigner-Eckart theorem

$$\langle t_1\,j_1\,m_1 | T_q^k | t_2\,j_2\,m_2 \rangle$$
$$= \langle t_1\,j_1 \| T^k \| t_2\,j_2 \rangle \langle j_2\,k\,m_2\,q \| j_1\,m_1 \rangle, \qquad (10)$$

where j_i and m_i again parametrize angular momentum and its z component, respectively, and t_i denotes some additional degeneracy. The right-hand side consists of the reduced matrix element, $\langle t_1\,j_1 \| T^k \| t_2\,j_2 \rangle$, which has no m dependence, and the Clebsch-Gordan coefficient (CGC), $\langle j_2\,k\,m_2\,q \| j_1\,m_1 \rangle$. If we substitute (j_3, m_3, t_3) for (k, q), P for T^k, and C for the CGC, we obtain the representation of an SU(2)-invariant rank-3 tensor,

$$T_{i_1 i_2 i_3} = T_{(j_1 m_1 t_1),(j_2 m_2 t_2),(j_3 m_3 t_3)}$$
$$= P_{(j_1 t_1),(j_2 t_2),(j_3 t_3)}\,C_{(j_1 m_1),(j_2 m_2),(j_3 m_3)}. \qquad (11)$$

The entire machinery of SU(2)-symmetric tensors is based on Eq. (11), as a tensor of arbitrary rank can now be constructed by multiplying rank-3 tensors and summing over common indices. Figure 5 illustrates the resulting compactification for a higher-dimensional tensor. In the upper-left corner, we start with a rank-4 tensor with unspecified directions of bonds, depicted as bidirectional arrows. The first step is to think of this object as a sequence of four rank-three tensors, leading to external indices, (1,2,3,4), given by the original tensor and internal indices, (0, 1, a, 4, 0). The directions of the latter are arbitrary; we set them here as going from left to right. We now apply the Wigner-Eckart theorem in the form of Eq. (11) to each of the four tensors, leading to the factorization in the second line of Fig. 5. The final step is to sum over common indices j_i, t_i, and m_i that do not connect P and C, which leads to the last line and is equivalent to the following formula for a rank-4 tensor:

$$T_{(j_1 m_1 t_1),(j_2 m_2 t_2),(j_3 m_3 t_3),(j_4 m_4 t_4)}$$
$$= \sum_{j_a} P_{(j_1 t_1),(j_2 t_2),(j_3 t_3),(j_4 t_4)}^{j_a} \cdot C_{(j_1 m_1),(j_2 m_2),(j_3 m_3),(j_4 m_4)}^{j_a}. \qquad (12)$$

The scheme depicted in Fig. 5 can easily be generalized to tensors of arbitrary rank r, leading to

$$T_{(j_1m_1t_1),\ldots,(j_rm_rt_r)}$$
$$= \sum_{j_{a_1},\ldots,j_{a_{r-3}}} P^{j_{a_1},\ldots,j_{a_{r-3}}}_{(j_1t_1),\ldots,(j_rt_r)} C^{j_{a_1},\ldots,j_{a_{r-3}}}_{(j_1m_1),\ldots,(j_rm_r)}. \quad (13)$$

In addition to r external indices, we thus obtain $r - 3$ independent, internal, indices connecting the reduced tensor elements with the CGCs. This degeneracy is a corollary of the triangle inequality of angular momenta, $j_1 - j_2 \leqslant j \leqslant j_1 + j_2$, and stands in contrast to a plain $U(1)$ symmetry, where Abelian quantum numbers just add up and lead to a Kronecker delta.

Adopting the terminology of Singh and Vidal [67], who performed the same decomposition using fusion-splitting trees, we term $C^{j_{a_1},\ldots,j_{a_{r-3}}}_{(j_1m_1),\ldots,(j_rm_r)}$ an *intertwiner*. Intertwiners inherit the orthogonality relations of CGCs in the form

$$\sum_{m_1,\ldots,m_r} C^{j_{a_1},\ldots,j_{a_{r-3}}}_{(j_1m_1),\ldots,(j_rm_r)} C^{j_{a'_1},\ldots,j_{a'_{r-3}}}_{(j_1m_1),\ldots,(j_rm_r)}$$
$$= N^{j_{a_1},\ldots,j_{a_{r-3}}}_{j_1,\ldots,j_r} \delta_{j_{a_1},j_{a'_1}}, \ldots, \delta_{j_{a_{r-3}},j_{a'_{r-3}}}, \quad (14)$$

where

$$N^{j_{a_1},\ldots,j_{a_{r-3}}}_{j_1,\ldots,j_r} = \sum_{m_1,\ldots,m_r} C^{j_{a_1},\ldots,j_{a_{r-3}}}_{(j_1m_1),\ldots,(j_rm_r)} C^{j_{a_1},\ldots,j_{a_{r-3}}}_{(j_1m_1),\ldots,(j_rm_r)}, \quad (15)$$

which can be used to flesh out the reduced tensors P for a given full tensor T,

$$P^{j_{a_1},\ldots,j_{a_{r-3}}}_{(j_1t_1),\ldots,(j_rt_r)} = \left(N^{j_{a_1},\ldots,j_{a_{r-3}}}_{j_1,\ldots,j_r}\right)^{-1}$$
$$\times \sum_{m_1,\ldots,m_r} C^{j_{a_1},\ldots,j_{a_{r-3}}}_{(j_1m_1),\ldots,(j_rm_r)} T_{(j_1m_1t_1),\ldots,(j_rm_rt_r)}. \quad (16)$$

To verify $SU(2)$ symmetry, one can then reinsert P into Eq. (13) and check if the initial and final T are equal.

The purpose of this scheme is thus to be able to operate on a compressed object P instead of the full tensor T, given that the model under investigation is $SU(2)$ invariant. Except for the initialization of the algorithm, the intertwiners C are not actually calculated, but serve as placeholding aids to derive the elementary tensor operations discussed in the following.
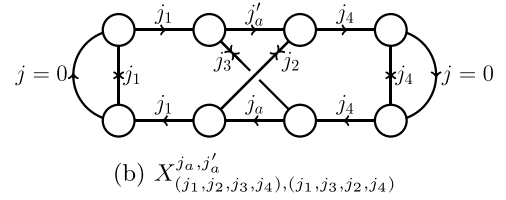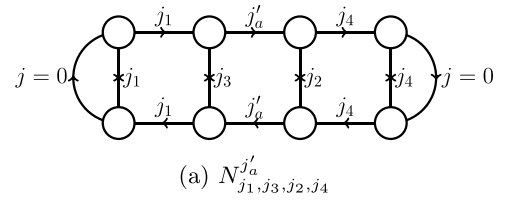
B. Permutation

Consider the permutation of two adjacent indices in a rank-4 tensor. Through Eq. (13), the two orders are defined as follows:

$$T_{(j_1m_1t_1),(j_2m_2t_2),(j_3m_3t_3),(j_4m_4t_4)}$$
$$= \sum_{j_a} P^{j_a}_{(j_1t_1),(j_2t_2),(j_3t_3),(j_4t_4)} \cdot C^{j_a}_{(j_1m_1),(j_2m_2),(j_3m_3),(j_4m_4)}, \quad (17)$$

$$T_{(j_1m_1t_1),(j_3m_3t_3),(j_2m_2t_2),(j_4m_4t_4)}$$
$$= \sum_{j'_a} P^{j'_a}_{(j_1t_1),(j_3t_3),(j_2t_2),(j_4t_4)} \cdot C^{j'_a}_{(j_1m_1),(j_3m_3),(j_2m_2),(j_4m_4)}. \quad (18)$$

Note how the first internal index j_a differs from the second j'_a. Since we want all orders of indices to be consistent with respect to the decomposition in Fig. 5, we think of a given

(a) $N^{j'_a}_{j_1,j_3,j_2,j_4}$

(b) $X^{j_a,j'_a}_{(j_1,j_2,j_3,j_4),(j_1,j_3,j_2,j_4)}$

(c) $\left(N^{j'_a}_{j_1,j_3,j_2,j_4}\right)^{-1} \cdot X^{j_a,j'_a}_{(j_1,j_2,j_3,j_4),(j_1,j_3,j_2,j_4)}$

FIG. 6. Permutation of an $SU(2)$-invariant rank-4 tensor.

$P^{j_a}_{(j_1t_1),(j_2t_2),(j_3t_3),(j_4t_4)}$ as being multiplied by the corresponding intertwiner, then permuted as usual, and finally compactified through Eq. (16), leading to the new reduced tensor $P^{j'_a}_{(j_1t_1),(j_3t_3),(j_2t_2),(j_4t_4)}$,

$$P^{j'_a}_{(j_1t_1),(j_3t_3),(j_2t_2),(j_4t_4)}$$
$$= \left(N^{j'_a}_{j_1,j_3,j_2,j_4}\right)^{-1} \sum_{j_a} P^{j_a}_{(j_1t_1),(j_2t_2),(j_3t_3),(j_4t_4)}$$
$$\cdot X^{j_a,j'_a}_{(j_1,j_2,j_3,j_4),(j_1,j_3,j_2,j_4)}, \quad (19)$$

where

$$X^{j_a,j'_a}_{(j_1,j_2,j_3,j_4),(j_1,j_3,j_2,j_4)}$$
$$= \sum_{m_1,m_2,m_3,m_4} C^{j_a}_{(j_1m_1),(j_2m_2),(j_3m_3),(j_4m_4)}$$
$$\cdot C^{j'_a}_{(j_1m_1),(j_3m_3),(j_2m_2),(j_4m_4)}. \quad (20)$$

Figures 6(a) and 6(b) depict the clusters of CGCs that make up N and X, respectively, where each circle represents a CGC and each line a common index. Since Eq. (19) requires only their ratio, the CGCs at the edges cancel out and lead to Fig. 6(c). The left cluster can easily be calculated using the orthogonality relations of CGCs. The right cluster is proportional to a Wigner 6-j symbol and can be calculated analytically using the Racah-formula or numerically by generating the CGCs and summing over indices according to the sketch. Note that for every single j_a, there are potentially multiple j'_a.

The best way to implement an arbitrary permutation is through successive permutations of two adjacent indices. These computations are the only expensive ones within the $SU(2)$ bookkeeping, which is why the resulting prefactors

(a) $N_{j_1, j_2, j_5, j_6}^{j_a}$

(b) $Y_{(j_1, j_2, j_3, j_4), (j_4, j_3, j_5, j_6), (j_1, j_2, j_5, j_6)}^{j_a}$

(c) $\left(N_{j_1, j_3, j_2, j_4}^{j_a} \right)^{-1} \cdot Y_{(j_1, j_2, j_3, j_4), (j_4, j_3, j_5, j_6), (j_1, j_2, j_5, j_6)}^{j_a}$
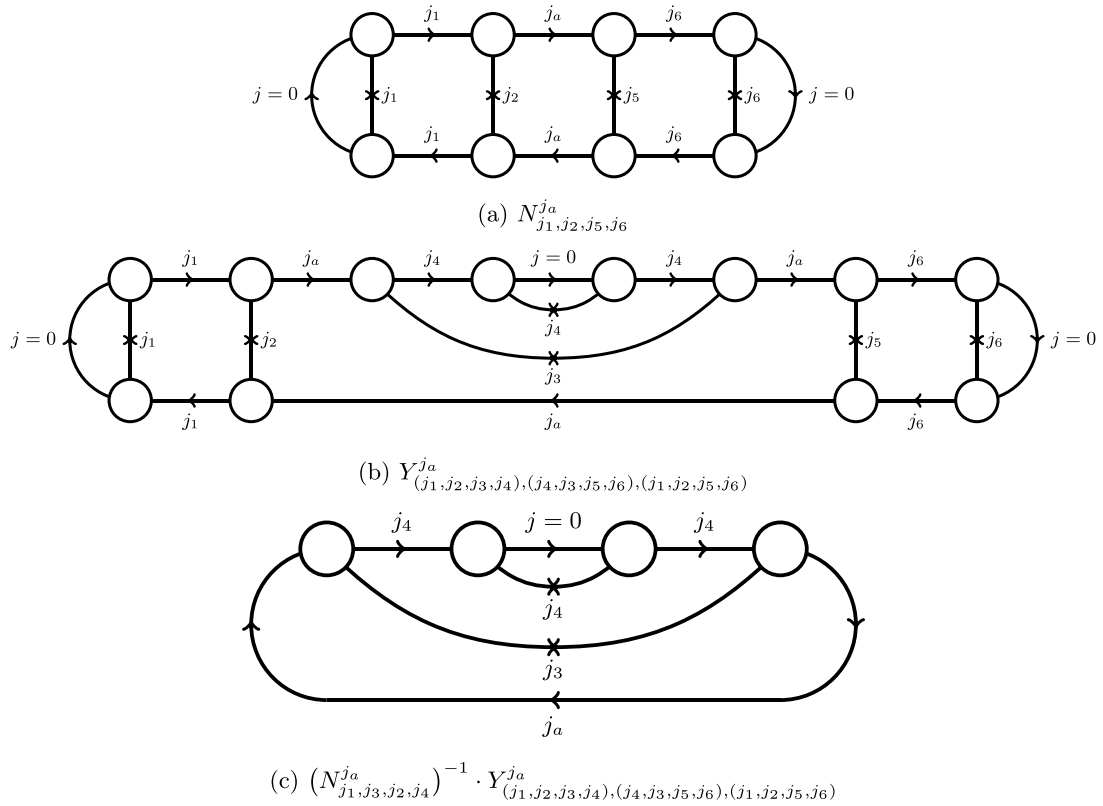
FIG. 7. Contraction of two SU(2)-invariant rank-4 tensors.

$N^{-1} \cdot X$ are best calculated on demand and then cached for future retrieval.

C. Contraction

As a minimal, but sufficiently general, example for the contraction of tensors, we consider two rank-4 tensors, T^1 and T^2, summed over two common indices. The first step is to permute both orders of indices using the scheme outlined in the previous section to bring them into the following form:

$$
\begin{aligned}
&T_{(j_1 m_1 t_1), (j_2 m_2 t_2), (j_3 m_3 t_3), (j_4 m_4 t_4)}^1 \\
&= \sum_{j_a} P_{(j_1 t_1), (j_2 t_2), (j_3 t_3), (j_4 t_4)}^{1, j_a} \cdot C_{(j_1 m_1), (j_2 m_2), (j_3 m_3), (j_4 m_4)}^{j_a},
\end{aligned}
$$

(21)

$$
\begin{aligned}
&T_{(j_3 m_3 t_3), (j_4 m_4 t_4), (j_5 m_5 t_5), (j_6 m_6 t_6)}^2 \\
&= \sum_{j_a} P_{(j_3 t_3), (j_4 t_4), (j_5 t_5), (j_6 t_6)}^{2, j_a} \cdot C_{(j_4 m_4), (j_3 m_3), (j_5 m_5), (j_6 m_6)}^{j_a}.
\end{aligned}
$$

(22)

Note the different orders of indices in Eq. (22). The indexing (3,4,5,6) on the degenerate level allows for the contraction to be implemented as a matrix-matrix multiplication, whereas the order (4,3,5,6) on the structural level leads to a trivial cluster of CGCs. Analogously to Sec. III B, we think of the reduced tensors P^1 and P^2 as being multiplied by their intertwiners, contracted, and then compactified via Eq. (16). If the

result is defined by

$$
\begin{aligned}
&T_{(j_1 m_1 t_1), (j_2 m_2 t_2), (j_5 m_5 t_5), (j_6 m_6 t_6)}^3 \\
&= \sum_{j_a} P_{(j_1 t_1), (j_2 t_2), (j_5 t_5), (j_6 t_6)}^{3, j_a} \cdot C_{(j_1 m_1), (j_2 m_2), (j_5 m_5), (j_6 m_6)}^{j_a},
\end{aligned}
$$

(23)

the reduced tensor P^3 can be written as

$$
\begin{aligned}
&P_{(j_1 t_1), (j_2 t_2), (j_5 t_5), (j_6 t_6)}^{3, j_a} \\
&= \left(N_{j_1, j_2, j_5, j_6}^{j_a} \right)^{-1} \sum_{j_3, t_3, j_4, t_4} P_{(j_1 t_1), (j_2 t_2), (j_3 t_3), (j_4 t_4)}^{1, j_a} \\
&\quad \cdot P_{(j_3 t_3), (j_4 t_4), (j_5 t_5), (j_6 t_6)}^{2, j_a} \cdot Y_{(j_1, j_2, j_3, j_4), (j_4, j_3, j_5, j_6), (j_1, j_2, j_5, j_6)}^{j_a},
\end{aligned}
$$

(24)

where

$$
\begin{aligned}
&Y_{(j_1, j_2, j_3, j_4), (j_4, j_3, j_5, j_6), (j_1, j_2, j_5, j_6)}^{j_a} \\
&= \sum_{m_1, \dots, m_6} C_{(j_1 m_1), (j_2 m_2), (j_3 m_3), (j_4 m_4)}^{j_a} \\
&\quad \cdot C_{(j_4 m_4), (j_3 m_3), (j_5 m_5), (j_6 m_6)}^{j_a} \cdot C_{(j_1 m_1), (j_2 m_2), (j_5 m_5), (j_6 m_6)}^{j_a}. \quad (25)
\end{aligned}
$$

The clusters N and Y are depicted graphically in Figs. 7(a) and 7(b), respectively. Due to the advantageous order of indices, the cluster Y and thus the ratio given in Fig. 7(c) is now a trivial sum over CGCs that can be processed using their orthogonality relations and the identity

$$
\langle j_1 \, m_1 \, j_2 \, m_2 || j \, m \rangle = (-1)^{j_1 + j_2 - j} \langle j_2 \, m_2 \, j_1 \, m_1 || j \, m \rangle. \quad (26)
$$

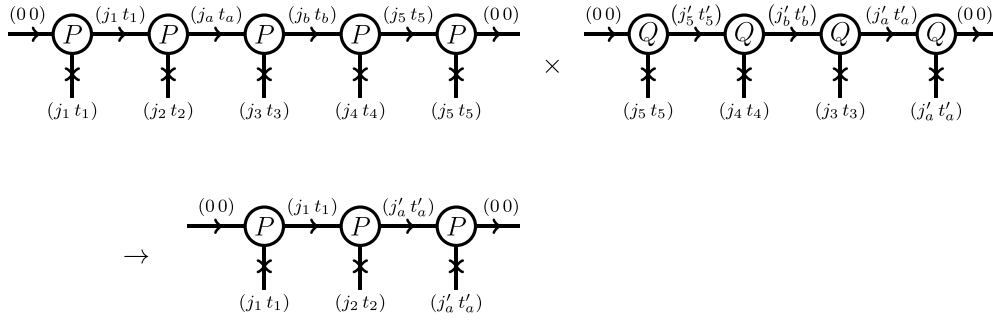The generalization to arbitrary contractions is straightforward.

FIG. 8. Fusion of three indices in an SU(2)-invariant rank-5 tensor.

The approach presented above can be shown to be equivalent to the use of X symbols introduced by Weichselbaum [69].

D. Charge fusion

In order to perform singular value decompositions in tensor networks, indices need to be fused beforehand and split afterwards. For this purpose, we take a rank-5 tensor as an example and conceptualize its reduced version as a sequence of rank-three tensors with internal angular momenta j_a and j_b and fictitious, internal, dense, indices t_a and t_b, as depicted in Fig. 8. Suppose that the indices 3, 4, and 5 have to be fused by an isometry Q. Due to the selection rules of CGCs, the external angular momenta, j_3, j_4, and j_5, do not uniquely identify the outcome of fusion, in contrast to the case of U(1) symmetry, for which the outcome would simply be the sum of Abelian quantum numbers. Instead, the final index of the isometry is dictated by the internal index of the reduced tensor, in this case j_a. To isolate it, we think of the reduced isometry as a sequence of rank-3 tensors as well, shown as four Q's with the reverse order of indices in Fig. 8. Adjacent P's and Q's are now contracted from the inside out. The first two yield a result proportional to $\delta_{j_5, j_5'} \delta_{t_5, t_5'}$ due to SU(2) invariance, which can be absorbed to either the left or the right. The next contraction eliminates $(j_b t_b)$, and the last one finally exposes $(j_a t_a)$. The dimension d_a of t_a is determined by

$$d_a = \sum_{j_3, j_4, j_5, j_b} (d_3 \, d_4 \, d_5)_{j_3, j_4, j_5, j_b}. \tag{27}$$

One may again contrast this with U(1) symmetry, where the sum over j_b would be missing. The next step is to determine the proper prefactor $Q^{j_b'}_{j_5\,j_4\,j_3\,j_a'}$ that yields unitarity. With this in mind, we note that the reduced isometry $Q^{j_b'}_{(j_5 t_5),(j_4 t_4),(j_3 t_3),(j_a' t_a')}$ is accompanied by the intertwiner $C^{j_b'}_{(j_5 m_5),(j_4 m_4),(j_3 m_3),(j_a' m_a')}$. To ensure $Q Q^\dagger = \mathbb{1}$, the following condition must then hold:

$$Q^{j_b'}_{j_5\,j_4\,j_3\,j_a'} Q^{j_b'}_{j_5\,j_4\,j_3\,j_a''} \cdot \sum_{m_5, m_4, m_3} C^{j_b'}_{(j_5 m_5),(j_4 m_4),(j_3 m_3),(j_a' m_a')}$$

$$\times C^{j_b'}_{(j_5 m_5),(j_4 m_4),(j_3 m_3),(j_a'' m_a'')} = \delta_{j_a', j_a''} \, \delta_{m_a', m_a''} \tag{28}$$

$$\Rightarrow Q^{j_b'}_{j_5\,j_4\,j_3\,j_a'} = \sqrt{\frac{2\,j_a' + 1}{N^{j_b}_{j_5, j_4, j_3, j_a'}}}. \tag{29}$$

As in the case of permutation and contraction, we assume that this example of charge fusion is sufficient to illustrate how to proceed in the general case.

E. Index reversal

The contraction of two tensors sometimes requires the reversal of indices, so that joint indices to be summed over are always ingoing for one and outgoing for the other. If all indices are initialized so that this condition is met in the beginning of a tensor network algorithm, and their directions are consistently reassigned during decompositions, it is sufficient to implement the reversal of all indices of a tensor only. To this end, we again consider our standard example of a rank-4 tensor, this time with specified directions of indices, say (out,in,in,out)

$$T_{(j_1 m_1 t_1),(j_2 m_2 t_2),(j_3 m_3 t_3),(j_4 m_4 t_4)}$$

$$= \sum_{j_a} P^{j_a}_{(j_1 t_1),(j_2 t_2),(j_3 t_3),(j_4 t_4)} \cdot C^{j_a}_{(j_1 m_1),(j_2 m_2),(j_3 m_3),(j_4 m_4)}. \tag{30}$$

We denote the flipped set of directions, (in,out,out,in), by a tilde

$$\tilde{T}_{(j_1 m_1 t_1),(j_2 m_2 t_2),(j_3 m_3 t_3),(j_4 m_4 t_4)}$$

$$= \sum_{j_a} \tilde{P}^{j_a}_{(j_1 t_1),(j_2 t_2),(j_3 t_3),(j_4 t_4)} \cdot \tilde{C}^{j_a}_{(j_1 m_1),(j_2 m_2),(j_3 m_3),(j_4 m_4)}. \tag{31}$$

Since the indices of a full tensor T can be flipped simultaneously without changing its value, $T = \tilde{T}$. Following the methods of the previous sections, we multiply a given P with its intertwiner C, flip the indices, and compactify using Eq. (16), applying the intertwiner $\tilde{C}$ with reversed external indices,

$$\tilde{P}^{j_a}_{(j_1 t_1),(j_2 t_2),(j_3 t_3),(j_4 t_4)} = \left(\tilde{N}^{j_a}_{j_1, j_2, j_3, j_4}\right)^{-1} \sum_{j_a} P^{j_a}_{(j_1 t_1),(j_2 t_2),(j_3 t_3),(j_4 t_4)}$$

$$\times Z^{j_a, j_a}_{(j_1, j_2, j_3, j_4),(j_1, j_2, j_3, j_4)}, \tag{32}$$

$$Z^{j_a, j_a}_{(j_1, j_2, j_3, j_4),(j_1, j_2, j_3, j_4)} = \sum_{m_1, m_2, m_3, m_4} C^{j_a}_{(j_1 m_1),(j_2 m_2),(j_3 m_3),(j_4 m_4)}$$

$$\cdot \tilde{C}^{j_a}_{(j_1 m_1),(j_2 m_2),(j_3 m_3),(j_4 m_4)}. \tag{33}$$

The resulting clusters N and Z are depicted in Figs. 9(a) and 9(b), respectively, and can again be processed by utilizing the orthogonality relations of CGCs.
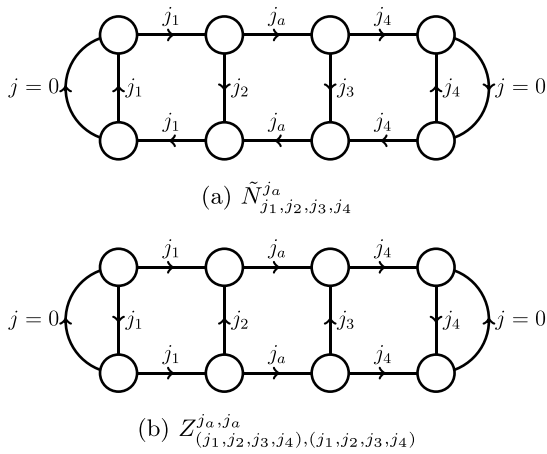
(a) $\tilde{N}^{j_a}_{j_1,j_2,j_3,j_4}$

(b) $Z^{j_a,j_a}_{(j_1,j_2,j_3,j_4),(j_1,j_2,j_3,j_4)}$

FIG. 9. Index reversal of an SU(2)-invariant rank-4 tensor.

IV. EVALUATING AND OPTIMIZING A PEPS

We now turn to extracting physical properties from an fPEPS (or, more precisely, a tensor network composing of fPEPS and PEPOs). First, we need to describe how to evaluate expectation values of fPEPS-PEPO tensor networks; they are required to implement optimization schemes as well as to measure observables. Subsequently, we describe two different schemes to variationally optimize fPEPS: local updates based on the iterative diagonalization of effective Hamiltonians for a single site or a single bond, and a global optimization of the fPEPS in which an approximation to the gradient of the energy functional is used as the basis of a gradient-based optimization scheme.

A. Expectation values

Given the representations of PEPSs and PEPOs described in previous sections, we now consider the calculation of

(a)

(b)

(c)

(d)

(e)

(f)
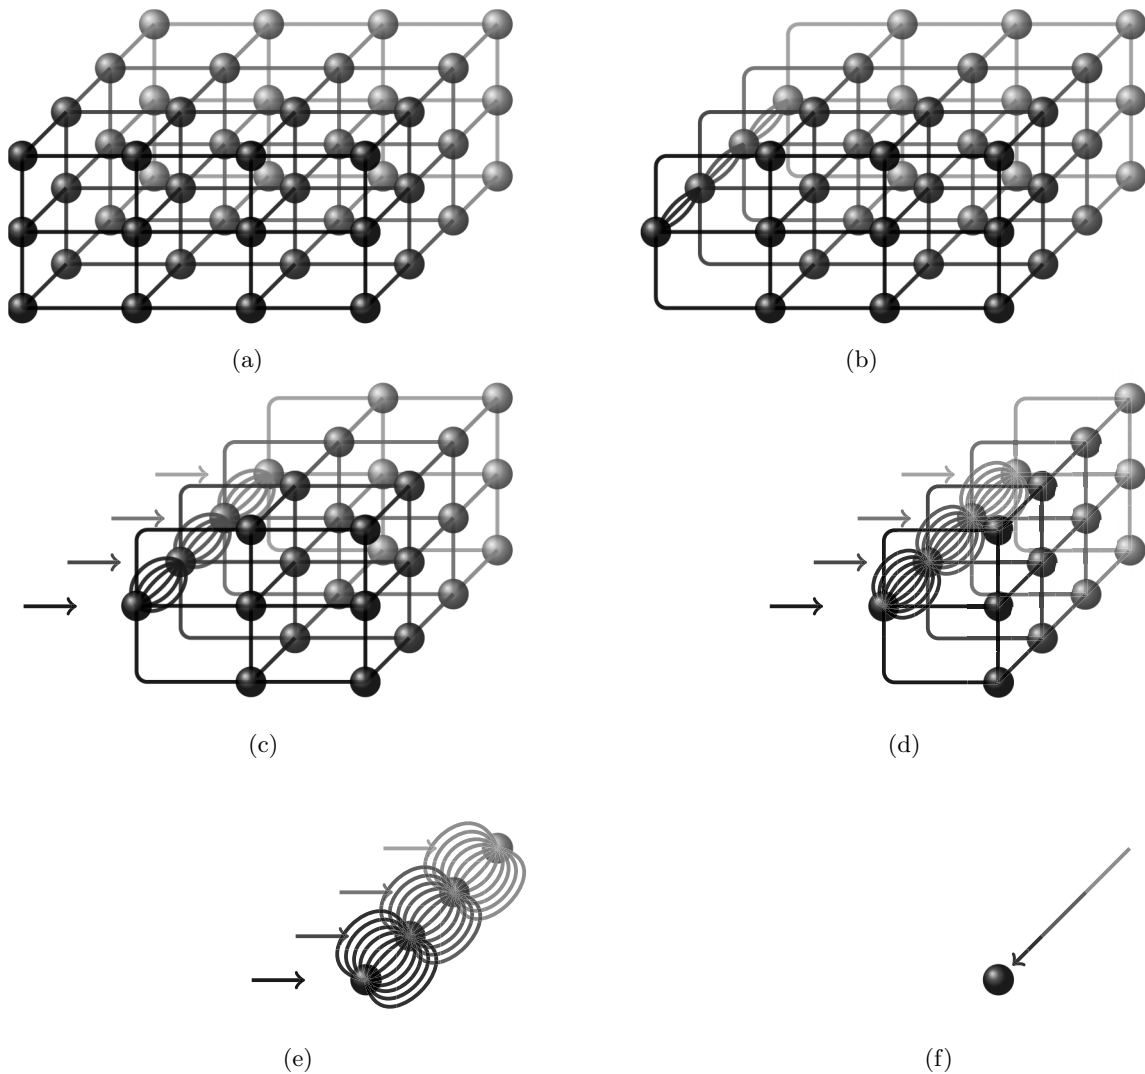
FIG. 10. Naive contraction of a PEPS-based tensor network. (a) $\langle \psi | \hat{O} | \psi \rangle$ as a tensor network, (b) Contraction of first environment, (c) Contraction of second environment, (d) Contraction of third environment, (e) Contraction of fourth environment, and (f) Contraction of remaining tensors.
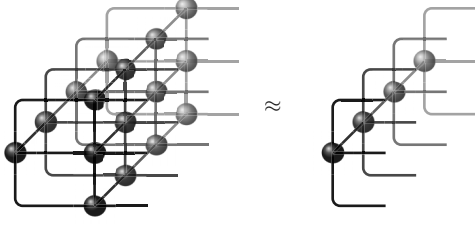
FIG. 11. Approximation of environment.

expectation values using these representations, i.e., $\langle \psi | \hat{O} | \psi \rangle$. The straightforward sequential contraction of an fPEPS-PEPO-fPEPS tensor network one might naively utilize to do this is illustrated in Fig. 10 for a 4×4 lattice. Figure 10(a) depicts the initial tensor network. The upper, middle, and lower planes are the fPEPS, PEPO, and adjoint fPEPS respectively. The subsequent contraction steps, depicted in Figs. 10(b)–10(e), successively build environments until the border is reached and the remaining tensors are contracted into a single number, yielding $\langle \psi | \hat{O} | \psi \rangle$ in Fig. 10(f).

The difficulty with this naive approach is evident: In every new contraction step, the rank of the intermediate tensors increases; this explosion of tensor complexity cannot be circumvented by choosing any other conceivable order of contractions. The exact contraction of fPEPS-based tensor networks thus scales exponentially with respect to system size and is therefore not feasible to carry out in realistic applications. However, we can circumvent this problem by constructing a sufficiently well-controlled approximation for each environment.

The environment approximation scheme we use here is depicted in Fig. 11; this scheme is is similar to the boundary-MPO construction used in Ref. [43]. Suppose we are given the arrangement of tensors on the left, which consist of environment blocks $E_{i-1,j}$, typically retained from a previous calculation step, fPEPS tensors $A_{i,j}$, their adjoints $A_{i,j}^\dagger$, and PEPO tensors $W_{i,j}$. The goal is to find the tensors $E_{i,j}$ on the right, which, as a whole, approximate the cluster on the left as well as possible by forming cumulative indices $\gamma_{i,j}$ with a predetermined maximum bond dimension χ [37]. If we interpret the former network as a full vector, i.e.,

$$|\psi\rangle = (E_{i-1,1} \cdot A_{i,1} \cdot W_{i,1} \cdot A_{i,1}^\dagger) \cdot$$
$$\ldots \cdot (E_{i-1,N} \cdot A_{i,N} \cdot W_{i,N} \cdot A_{i,N}^\dagger), \quad (34)$$

and the latter as a truncated vector,

$$|\tilde{\psi}\rangle = E_{i,1} \cdot E_{i,2} \cdot \ldots \cdot E_{i,N}, \quad (35)$$

the problem can be stated as finding the maximum of the fidelity

$$F = \frac{\langle \psi | \tilde{\psi} \rangle \langle \tilde{\psi} | \psi \rangle}{\langle \tilde{\psi} | \tilde{\psi} \rangle \langle \psi | \psi \rangle}, \quad (36)$$

whose solution, in general, is given by

$$\langle \tilde{\psi} | \tilde{\psi} \rangle = \langle \tilde{\psi} | \psi \rangle. \quad (37)$$

Using Eq. (37) to find the optimal $|\tilde{\psi}\rangle$ as a whole is not possible because $|\tilde{\psi}\rangle$ consists of multiple tensors, and contracting all of them is actually what we want to avoid. Instead,
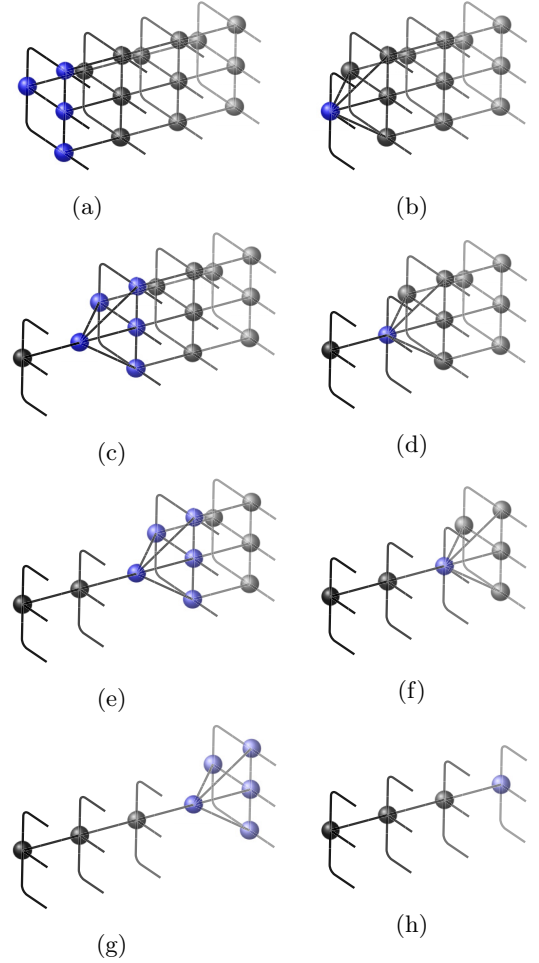

FIG. 12. Initialization of truncated environment $|\tilde{\psi}\rangle$. (a) Initial arrangement, (b) Contraction, (c) Decomposition, (d) Contraction, (e) Decomposition, (f) Contraction, (g) Decomposition, and (h) Contraction.

we start with a trial vector constructed as shown in Fig. 12. The cluster in Fig. 12(a) is the original $|\psi\rangle$, while the subsequent figures illustrate how to systematically forge a truncated environment with a bond dimension of one by alternately carrying out contractions and truncated SVDs. The outcome, Fig. 12(h), is then the starting point for an iterative algorithm.

The optimization of $|\tilde{\psi}\rangle$ proceeds similarly to that carried out to find the ground state in the density matrix renormalization group (DMRG): Most of its tensors are fixed in that $\langle \tilde{\psi} | \tilde{\psi} \rangle$ and $\langle \tilde{\psi} | \psi \rangle$ are calculated partially starting from both sides. The individual steps are shown in Fig. 13, which depicts an aerial view; here the tensors F and G represent intermediate contractions for $\langle \tilde{\psi} | \tilde{\psi} \rangle$ and $\langle \tilde{\psi} | \psi \rangle$, respectively. Only two sites, labeled j and $j+1$ here, are not contracted over, which makes it possible to determine the optimal joint environment block $X_{j,j+2} = E_{j,j+1} \cdot E_{j+1,j+2}$ by requiring that

$$\langle \tilde{\psi} | \tilde{\psi} \rangle / X_{j,j+2}^\dagger = \langle \tilde{\psi} | \psi \rangle / X_{j,j+2}^\dagger \quad (38)$$

in this reduced vector space. The calculation necessary to obtain the right-hand side is depicted in Fig. 14, yielding the inhomogeneity B. The contraction is best performed by calculating the left and right halves of the cluster separately, then
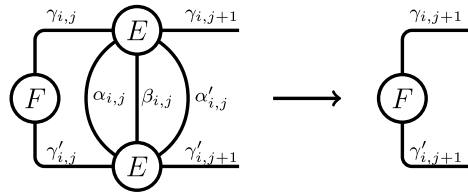
(a)
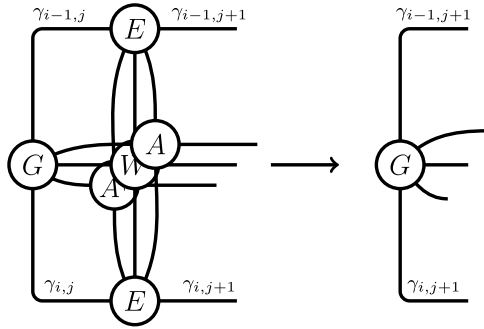
(b)

FIG. 13. Intermediate contraction steps. (a) Intermediate contraction steps of $\langle \tilde{\psi} | \tilde{\psi} \rangle$ and (b) Intermediate contraction steps of $\langle \tilde{\psi} | \psi \rangle$.

multiplying the two results with each other. Figure 15 shows how $F_{i,j}$ and $F_{i,j+2}$ are then contracted, yielding the rank-4 tensor M. Finding the solution to Eq. (38) now amounts to solving $M \cdot X = B$ for all columns of B, which is depicted in Fig. 16. However, if one performs a QR decomposition on the environment tensors E, they become isometries, and all F and thus M are identity matrices. The solution X is then simply the inhomogeneity B. Its decomposition into $E_{i,j}$ and $E_{i,j+1}$ is
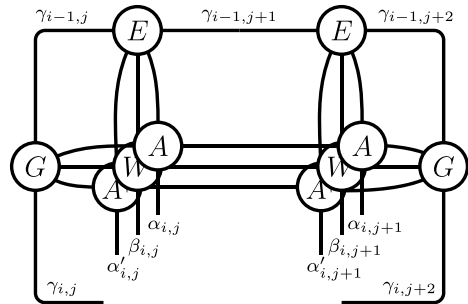
(a)
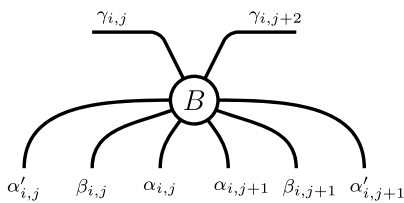
(b)

FIG. 14. Calculation of the inhomogeneity B. (a) Initial setup and (b) Contraction result.
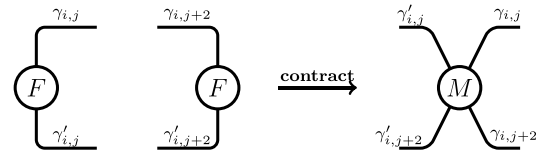
FIG. 15. Calculation of M.

depicted in Fig. 17. This generates the composite index $\gamma_{i,j+1}$, which has the predetermined maximum bond dimension χ, i.e., the maximum bond dimension of environment tensors. It is important not to confuse χ with the maximum number of virtual states within the PEPS, D. Note that algorithms based on MPSs as well as other acyclic tensor networks do not, in general, have a numerical parameter χ, as the corresponding environments are processed exactly.

The scheme presented above still contains one major drawback: If the maximum bond dimension is χ, the contraction in Fig. 14 generates a temporary number of states $\tilde{\chi} \gg \chi$, most of which are dropped in the subsequent truncation in Fig. 17 and are therefore superfluous. This issue is not restricted to finding the best virtual basis within environment tensors in our fPEPS-PEPO-scheme, but is a general computational bottleneck in tensor network algorithms, in particular, in the typical scheme in which the bond dimension between two adjacent tensors is significantly enlarged, then optimized via the maximum overlap or minimal energy, and subsequently truncated again. In order to reduce the cost of such steps, Gleis *et al.* [58] have recently developed the so-called "controlled bond expansion" (CBE) for calculating ground-state MPSs. In this method, one picks ξ orthogonal states with the largest weight, adds them to the current χ states in tangent space, optimizes the combined set of states, and, finally, truncates it to the desired dimension. If ξ is chosen such that $\tilde{\chi} \gg \chi + \xi$, one can thus perform a two-site optimization at one-site cost. In the following, we will sketch how to use their method for the environment approximation described above.

We start by formulating the completeness relations in Figs. 18(a) and 18(b) for the tensors $E_{i,j}$ and $E_{i,j+1}$, respectively. Due to the QR decomposition (described above), the approximated environment $| \tilde{\psi} \rangle$ is in a form that is essentially the mixed canonical form typically used for MPSs, and the rank-8 tensors $E \cdot E^{\dagger}$ on the left are projectors onto E, i.e., onto the tangent space. The right-hand sides of both figures consist of unit tensors, which are formally constructed as the tensor product of Kronecker deltas and are simply depicted as lines. These configurations define the orthogonal projectors $E' \cdot (E')^{\dagger}$, which are then used to enlarge the bond between two tensors in a controlled manner. Figure 19(a) shows how the contraction of the left orthogonal projector with the left half of Fig. 14(a) is carried out. The first step is to decompose the tensors G and E in the manner depicted in Fig. 19(b), while the new index $\delta_{i,j}$ is truncated from χ to ξ. Afterwards,
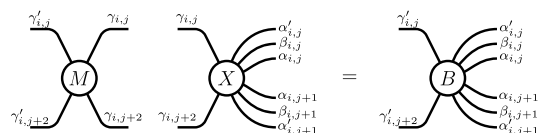
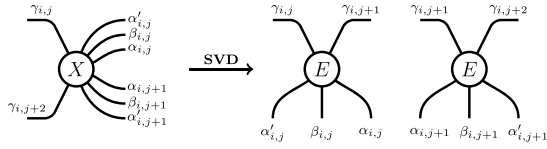

FIG. 16. Solving for a joint environment tensor.

FIG. 17. Decomposition of joint environment tensor.

both clusters are contracted and subtracted, yielding the left orthogonal block depicted in Fig. 19(c). A second truncated decomposition, Fig. 19(d), isolates the weight ΣV^T from the left block, which is then passed on to the right orthogonal block, Fig. 20(a). After the weight is absorbed, Fig. 20(b), the clusters are contracted and subtracted, as depicted in Fig. 20(c). The final decomposition, Fig. 20(d), yields the truncated complement V^T, which encodes those bonds orthogonal to the current basis with the largest weight [58]. If the current dimension of $\gamma_{i,j+1}$ between $E_{i,j}$ [Fig. 18(a)] and $E_{i,j+1}$ [Fig. 18(a)] is given by $\dim(\gamma_{i,j+1})$, the bond dimension of $\gamma'_{i,j+1}$ in Fig. 20(d) is set to

$$\dim(\gamma'_{i,j+1}) = \chi + \xi - \dim(\gamma_{i,j+1}). \tag{39}$$

In the final step of the CBE, we add V^T to $E_{i,j+1}$; here Eq. (39) ensures that the dimension of the bond in the sum does not exceed $\chi + \xi$. An additional QR decomposition and the contraction depicted in Fig. 13(b) allows us to solve for $E_{i,j}$ using the cluster depicted in Fig. 21. If the dimension of $\gamma_{i,j+1}$ is greater than χ, the sequence of operations depicted in Fig. 22 can be used to truncate the dimension back to χ.

To summarize, the CBE allows one to circumvent the operations depicted in Figs. 14 and 17 by instead using orthogonal projectors and a series of contractions and truncated SVDs,
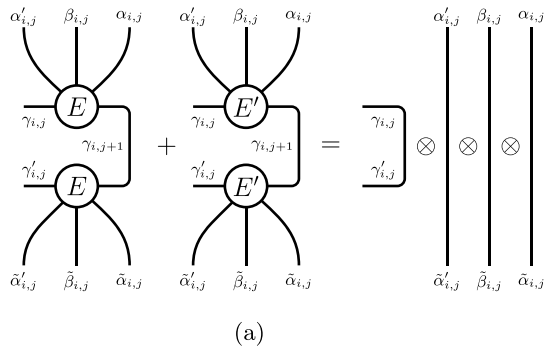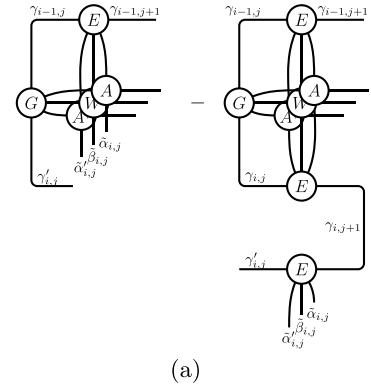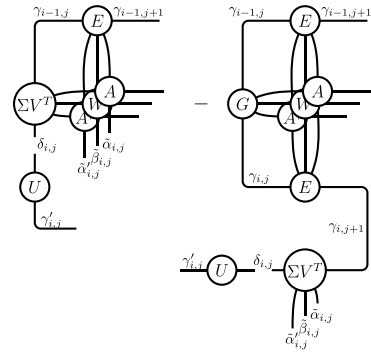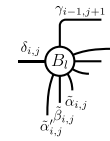

(a)

(b)

FIG. 18. Completeness relations of environment tensors. (a) Left completeness and (b) Right completeness.
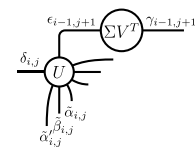

(a)

(b)

(c)

(d)

FIG. 19. Calculation of left orthogonal block. (a) Initial setup, (b) Truncated decompositions, (c) Contraction and subtraction, and (d) Truncated decomposition.

avoiding operating on a rank-8 tensor with indices of full dimension.

In order to progressively improve the environment, one can sweep repeatedly from left to right and from right to left in a manner similar to the optimization of an MPS in the DMRG. The state $|\tilde{\psi}\rangle$ then converges to the best approximation of $|\psi\rangle$ within the given maximum bond dimension χ. Thus, the scaling of the computational cost of contracting PEPS-based tensor networks with system size becomes linear rather than exponential if the contraction is carried out approximately in this manner.
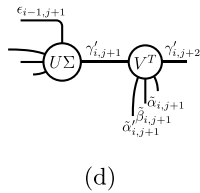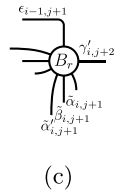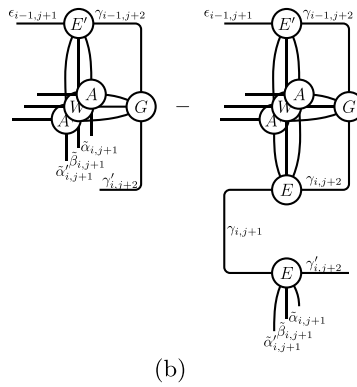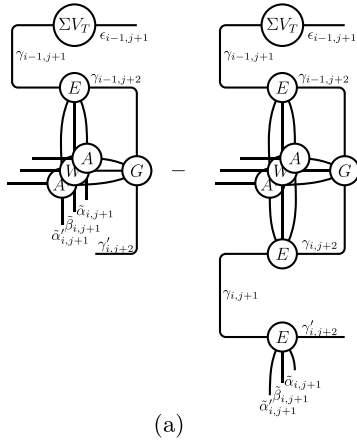
(a)

(b)

(c)

(d)

FIG. 20. Calculation of right orthogonal block. (a) Initial setup, (b) Contraction of weight from left block, (c) Contraction and subtraction, and (d) Truncated decomposition.
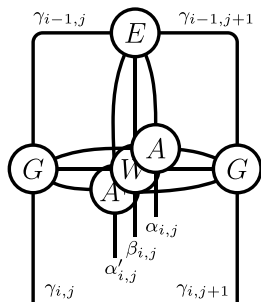

FIG. 21. Calculation of left environment tensor.
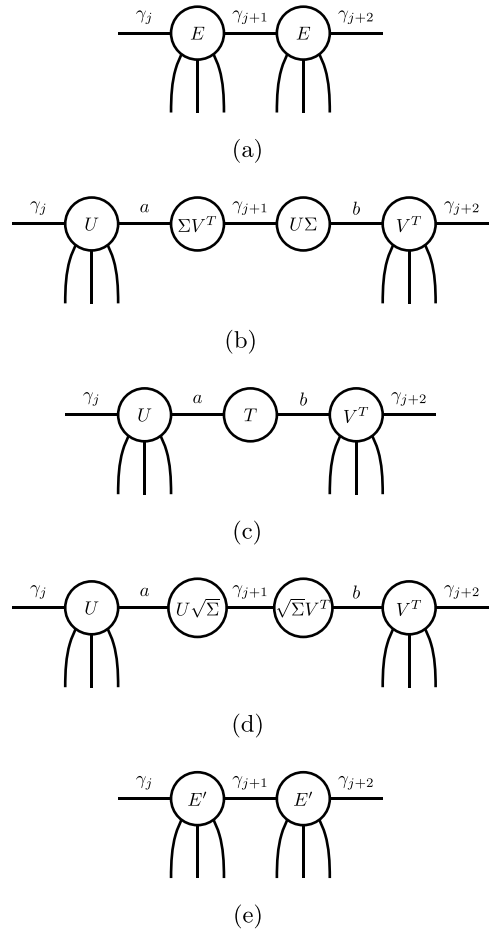

(a)

(b)

(c)

(d)

(e)

FIG. 22. Truncation of two adjacent environment tensors. (a) Initial setup, (b) Decompositions, (c) Contraction, (d) Truncated decomposition, and (e) Contractions.

Note that the approach presented above is not the only possibility for approximating environments. The original algorithm upon which our scheme is based, see Refs. [37,70], optimizes the environment tensors individually in a single-site manner. While this approach can be used in addition to the CBE optimization described above, on its own, it does not change the basis of the environment from that selected by the initialization and thus cannot converge to the best solution, as bonds are not optimized. In another approximation introduced in Ref. [71], physical indices rather than virtual indexes are bundled, which turns out to be less accurate than the approach of Ref. [43]. The least expensive way to approximate the environment is via a so-called "simple update" [72], which depends solely on the SVD of two adjacent environment tensors; however, this is also insufficiently accurate. Lubasch *et al.* [43] have proposed a compromise between the simple update and a full sweep by considering a cluster around the environment tensor one wants to optimize. This approach may be explored in addition to the scheme presented above, but introduces the cluster size as an additional numerical parameter. In addition, it is unclear to what extent this approximation would be beneficial for evaluating the PEPS-PEPO-network that we treat here.

B. Variational optimization

Having now covered all the necessary preliminaries, we discuss in the following two different methods for variationally optimizing an fPEPS to approximate the ground state. The first method, which we will term *local updates* involves contracting the fPEPS-PEPO-fPEPS until a local effective Hamiltonian, either for one site or for two sites (i.e., a bond), is formed. We diagonalize this Hamiltonian, which satisfies a generalized eigenvalue equation, using an iterative diagonalization scheme and then optimize the fPEPS locally, keeping the size of the basis fixed. This local-update optimization method is in the spirit of DMRG optimization of an MPS using the iterative diagonalization of an effective Hamiltonian followed by singular value decomposition. In the second method, termed *gradient updates*, we keep the basis fixed and optimize all local tensors in the fPEPS simultaneously using a gradient-based nonlocal optimization method.

Both optimization methods have shortcomings. In particular, local bond updates temporarily expand the Hilbert space and thus form the optimal virtual basis in a DMRG-like fashion, but are limited by the property that only two adjacent PEPS tensors are optimized at each step. Gradient updates are complementary to local updates in that all PEPS tensors are varied simultaneously to minimize the energy; however, this is done while keeping the virtual basis fixed.

We start with a rectangular lattice of a given width and height, the corresponding representation of the Hubbard Hamiltonian as a PEPO, and a particular sector of quantum numbers within which we want to find the ground state. Here we assume that we initialize the algorithm by converting a product state within that sector into an fPEPS. (Other initializations, such as a suitable randomized state, are also possible.)

1. Local updates

Let us consider first the case of two-site local optimization: assume that we want to optimize two adjacent bulk tensors, say $A_{x,y}$ and $A_{x+1,y}$. In order to do this, we first partially contract the network of the Hamiltonian, $\langle \text{fPEPS}|\text{PEPO}|\text{fPEPS}\rangle$, as depicted in Fig. 10(a), and also the network of the norm, $\langle \text{fPEPS}|\text{fPEPS}\rangle$. We do this by building environments from above and below using the scheme described in Sec. IV A. This leaves us with four tensor sequences, two for the energy, $E_{i,y+1}$ and $E_{i,y-1}$, and two for the norm, $N_{i,y+1}$ and $N_{i,y-1}$. Subsequently, we calculate environment blocks within row y by starting at both $x = 0$ and $x = x_{\max}$ and contracting according to the scheme depicted in Fig. 23. The tensors $A_{i,y}$, $W_{i,y}$, and $A^\dagger$ belong to the fPEPS, PEPO, and adjoint fPEPS, respectively; note that the latter two are partially obscured in Fig. 23. In the left diagram, $E_{i-1,y}$ does not exist in the initial step at the edge of the lattice but, in subsequent steps, contains the outcome of the previous contraction for $i > 0$. The contractions on the right take place analogously. One can picture this process as consisting of two zippers that close in on the two sites (x, y) and $(x + 1, y)$ from either side.

The configuration at the end of this process is depicted in Fig. 24. The effective Hamiltonian H_{eff}, Fig. 24(a), consists of six environment tensors E as well as the last two tensors of the PEPO that have not been contracted, $W_{x,y}$ and $W_{x+1,y}$, whereas
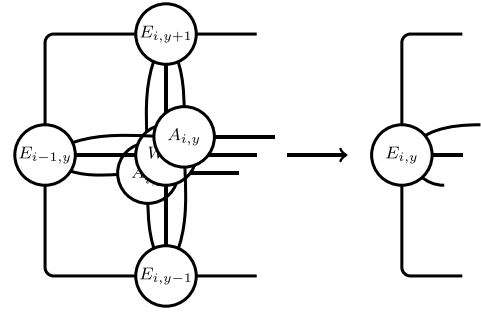

FIG. 23. Zipper contraction of blocks between two approximated environments.

the environment for the norm N_{env}, Fig. 24(b), consists of six norm tensors N only. Bonds connecting two E's or two N's are cumulative and are created during the approximation of the environment. Those connecting both W's or one W and one E represent states in an FSM. The free indices protruding from E's or N's encode quantum-mechanical entanglement within the PEPS, and the dangling links of both W's are physical Hilbert spaces. The clusters H_{eff} and N_{env} satisfy the generalized eigenvalue problem

$$H_{\text{eff}} |\psi\rangle = E N_{\text{env}} |\psi\rangle. \tag{40}$$

Its solution is the wave function $|\psi\rangle = A_{x,y} \cdot A_{x+1,y}$ that minimizes the energy E. Note that PEPSs cannot be put into any equivalent of the Schmidt form, in which the norm of a single tensor is equal to the norm of the overall wave function. If this were the case, we would have been able to find a gauge so that $\langle \psi | N_{\text{env}} | \psi \rangle = \langle \psi || \psi \rangle$, as in the DMRG.
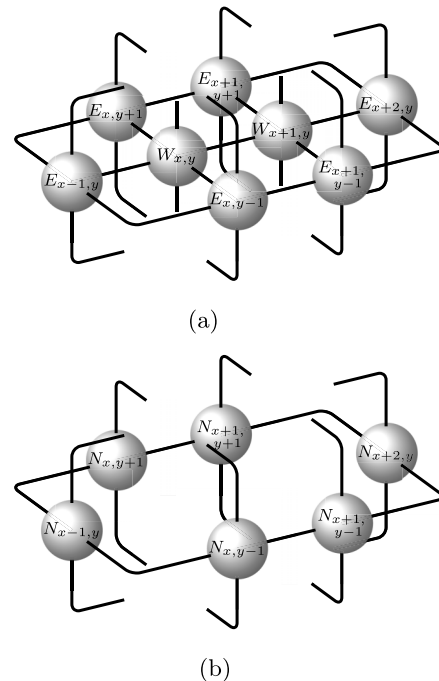

(a)

(b)

FIG. 24. Two-site optimization. (a) Effective Hamiltonian H_{eff} and (b) Norm environment N_{env}.
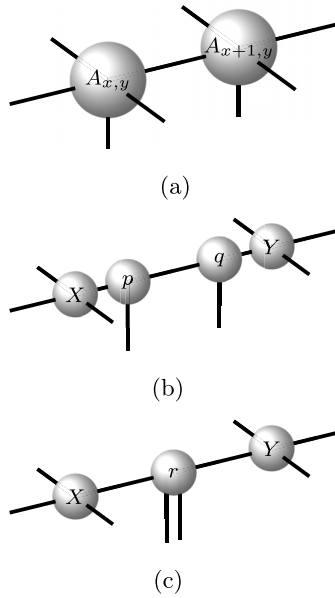
FIG. 25. Decomposition of PEPS tensors. (a) Two adjacent PEPS tensors, (b) Decompositions, and (c) Contraction.

In principle, one could now proceed with an optimization of $|\psi\rangle$ by carrying out iterative diagonalization directly. However, it is more efficient to first employ the bond optimization proposed by Corboz [55], whose initialization is illustrated in Fig. 25. We decompose the PEPS tensors $A_{x,y}$ and $A_{x+1,y}$, Fig. 25(a), into (X, p) and (q, Y), respectively, Fig. 25(b). We then contract the middle tensors, p and q, to form r, which carries both of the physical indices, Fig. 25(c). The tensors X and Y are then absorbed into the environment, leading to the decompositions $\hat{H} = H_{\text{eff}} \cdot X \cdot X^\dagger \cdot Y \cdot Y^\dagger$ and $\hat{N} = N_{\text{env}} \cdot X \cdot X^\dagger \cdot Y \cdot Y^\dagger$, which are depicted in Figs. 26(a) and 26(b), respectively. After applying the gauge proposed by Lubasch *et al.* [44], the tensor r can be optimized using a variation of the Davidson algorithm [73] adapted for the generalized eigenvalue problem. Subsequently, we factorize
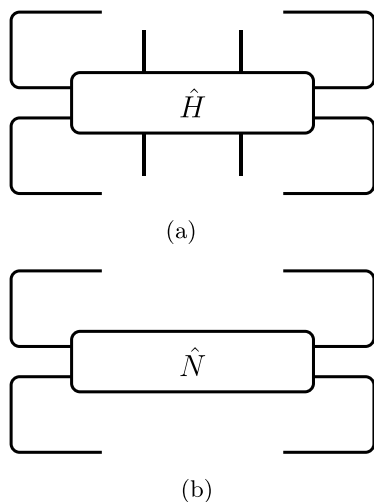

FIG. 26. Bond optimization. (a) Effective Hamiltonian $\hat{H}$ and (b) Norm environment $\hat{N}$.
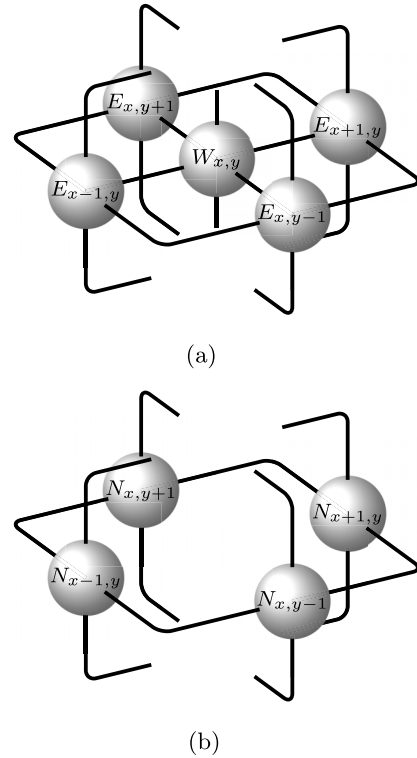

FIG. 27. One-site optimization. (a) Effective Hamiltonian and (b) Norm environment.

r back into p and q and then truncate the bond in between to a predetermined maximum number of states D. We then carry out a full update [46,74], which strips the wave function of cyclic entanglement [75] and improves the convergence.

In addition to bond optimization, one can also carry out one-site optimization, in which we again start with the configuration of Fig. 24, then apply the contraction of Fig. 23 to one open site, and, finally, arrive at the setup depicted in Fig. 27 for (a) the effective Hamiltonian and (b) the norm environment. Note that bond optimizations are able to explore new subsectors of the Hilbert space, analogously to a two-site DMRG algorithm for MPSs, whereas one-site optimizations can only optimize within a given basis.

Bond optimization is beset by the same problem as the environment approximation described in Sec. IV A: If the maximum bond dimension is D, the effective Hamiltonian of Fig. 26(a) will generate $\tilde{D} \gg D$ states, most of which will be dropped in the subsequent truncation. To avoid this, we again apply the CBE to circumvent contracting large tensors so that we can perform the bond optimization at one-site cost. Here we will not reprise the details of the CBE, but will instead only describe the adaptations necessary to apply it to PEPS tensors. First, we have to reexamine the bond between two tensors, as shown in Fig. 28. The product $r = A \cdot \sigma \cdot B$ is what we would normally optimize as a whole within the bond optimization. Note that, for MPSs, the norm environment N factorizes into two terms. We now choose the bond matrix σ so that it satisfies the weighted trace gauge [75], which can be seen as the generalization of left- and right normalization for cyclic tensor networks. The resulting orthogonality conditions are depicted in Fig. 29 for (a) the left norm and (b) the right norm.
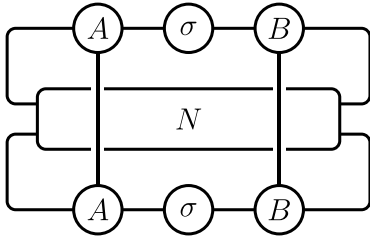
FIG. 28. Bond setup.

Subsequently, the completeness relations, depicted in Fig. 30, can be used to flesh out the orthogonal projectors $A' \cdot (A')^{\dagger}$ and $B' \cdot (B')^{\dagger}$. Aside from these two customizations, the CBE for PEPS tensors is structurally identical to the original version developed for MPSs [58] and the adaptation for environment tensors described in Sec. IV A.

The overall procedure for variational optimization via local updates for a 4×4 lattice is illustrated in Fig. 31. The sketches are to be read row-wise and from left to right, as indicated by the numbering. Sketch 1 in the upper left corner is a reduced aerial projection of Fig. 10(a) and depicts the starting point. The red arrows in sketch 2 denote three successive constructions of approximate environments for both the energy and the norm, which isolates the lowest row. The blue arrow in sketch 3 represents the zipper contraction of the two sites on the right, which prepares the system for the two-site optimization of the lower-left corner using the scheme described above. Subsequently, we make two steps to the right, optimize both times via bond optimization or CBE, reverse direction, and again optimize twice. In sketch 8, the treatment of the next row is prepared by building the environment of the recently processed first row. Sketches 9 to 13 illustrate the optimization of row 2 in the same order. The remaining sketches follow the same pattern, repeating until the uppermost row is reached. As a whole, the 25 sketches illustrate one horizontal sweep, which contains four DMRG-like optimizations of each row and, in the case of the Hubbard model, moves charge and spin across the lattice to minimize the energy. A vertical sweep can be implemented simply by rotating every sketch in Fig. 31 by 90 degrees. We denote a horizontal sweep followed by a vertical sweep a "full sweep".
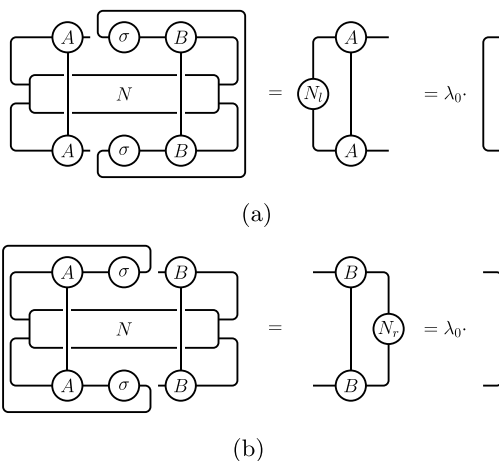

(a)

(b)

FIG. 29. Weighted trace gauge. (a) Left norm and (b) Right norm.
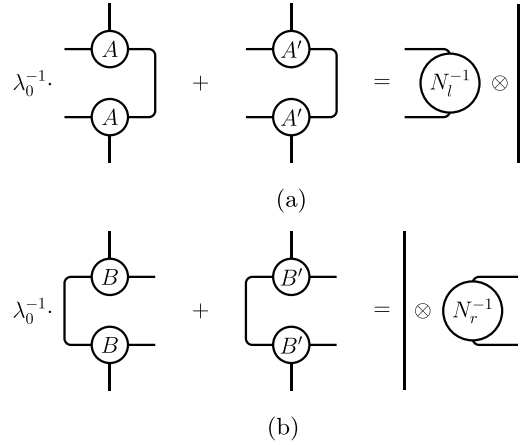
FIG. 30. Completeness relations. (a) Left completeness and (b) Right completeness.

As a final remark, we note that, due to the approximate construction of environments, the Rayleigh-Ritz variational principle is, unfortunately, violated. Therefore, the energy of the PEPS is not necessarily bounded from below by the ground-state energy of the Hamiltonian.

2. Gradient updates

Assuming that we have an fPEPS within a particular, fixed, virtual basis, we now consider the functional of the energy, varying one particular site i,

$$f(A_i) = \frac{A_i H_i A_i}{A_i N_i A_i} = \frac{e}{n}. \tag{41}$$

Here the A_i are elements of a PEPS tensor reshaped to a vector, while H_i and N_i are the effective Hamiltonian [Fig. 27(a)] and norm environments [Fig. 27(b)], respectively, reshaped to matrices. Note that H_i and N_i depend on all $A_{j \neq i}$. The gradient

FIG. 31. Horizontal sweep in variational optimization of fPEPS on a 4×4 lattice.

with respect to one site reads

$$\nabla f(A_i) = \frac{2}{n^2}(n H_i - e N_i)A_i. \qquad (42)$$

The complete gradient can be constructed from one-site gradients by concatenation,

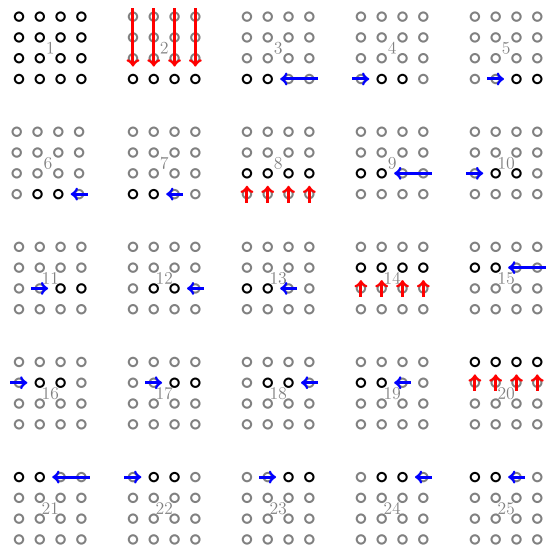$$\nabla f = (\nabla f(A_1), \nabla f(A_2), \ldots, \nabla f(A_N)). \qquad (43)$$

Once f and ∇f are available, one can employ any suitable gradient-based optimization algorithms to optimize the set of local tensors A_i. In the context of tensor networks, the conjugate gradient method was used by Vanderstraeten *et al.* [50] to carry out such an optimization within iPEPS. Here we instead utilize the L-BFGS algorithm [76,77], which is a memory-efficient, quasi-Newton method. We find the L-BFGS method to be substantially more effective, as might be expected since it is a second-order method, while conjugate gradient is first order.

V. RESULTS

In this section, we apply the fPEPS methods described above to the two-dimensional Hubbard model with Hamiltonian (1) and open boundary conditions on lattices ranging from 3×3 to 8×8. Our goal is to test the efficacy and performance of the method and to investigate its convergence as a function of the algorithmic parameters: the bond dimension D, the environment dimension χ, and the nature and number of update steps, where both local updates (Sec. IV B 1) and a combination of local and gradient updates (Sec. IV B 2) will be carried out. As a measure of convergence, we use primarily the ground-state energy, which we compare to (numerically) exact or, for larger lattice sizes, accurate variational results. As an accurate variational estimate of the ground-state energy E_0, we use a highly efficient DMRG program developed by G. Ehlers [78] keeping 4000 states for all simulations. This estimate is numerically exact for 3×3 and 4×4 lattices, should be quite accurate for 6×6 lattices, and will give a fairly accurate upper bound to the ground-state energy for 8×8 lattices. In addition to ground-state energies, we will also investigate the behavior of local observables, in particular, local hole density $1 - \langle n_i \rangle$, with $n_i = n_{i,\uparrow} + n_{i,\downarrow}$ and local spin density

$$\langle \mathbf{S}_i^2 \rangle = \tfrac{3}{4}\langle n_i - 2n_{i,\uparrow}n_{i,\downarrow} \rangle.$$

For the case of U(1) symmetry, we will also calculate $\langle S_i^z \rangle = \langle \tfrac{1}{2}(n_{i,\uparrow} - n_{i,\downarrow}) \rangle$, which is identically zero for an SU(2)-symmetric state, but can be nonzero for a U(1)-symmetric state, as an approximate numerical algorithm such as fPEPS can numerically break spin-inversion symmetry.

Here we treat the ground state in zero magnetic field only, so that both the total spin S and its z component S_z are taken to be zero for all calculations. This setup leads to two possible choices of symmetry groups: The first is $U(1)_{\text{spin}} \otimes U(1)_{\text{charge}}$, describing the local conservation of both the deviation from half-filling c_z, and the z component of the spin s_z. (We use small letters to denote conserved quantum numbers on sites, virtual indices, and state kets and capital letters to denote the corresponding conserved quantum numbers for states of the entire lattice.) The corresponding bonds of tensors are then

parameterized by $|s_z, c_z, t\rangle$, with t iterating over any additional degeneracy. within a symmetry sector of given c_z and s_z.

The second choice is to take advantage of the full spin-rotation symmetry and classify all states using $SU(2)_{\text{spin}} \otimes U(1)_{\text{charge}}$. This yields the far richer states $|s, c_z, t\rangle$ that, for a given total spin s, encompass the entire spin multiplet, that is, the set of states with $s_z = -s, \ldots, s_z = s$. In the following, we designate the use of the $U(1)_{\text{spin}} \otimes U(1)_{\text{charge}}$ symmetry as "U(1)" and the use of the $SU(2)_{\text{spin}} \otimes U(1)_{\text{charge}}$ symmetry as "SU(2)".

The central parameter for controlling the variational accuracy in fPEPS simulations (and in tensor network algorithms in general) is the number of virtual states D, usually referred to as the "bond dimension". Since we have to approximate the contraction of the tensor network for any optimization scheme, we must also choose the dimension χ of the environment tensors, as defined in Sec. IV A for a given set of model parameters, a given symmetry group, and a given bond dimension D. We do this empirically by choosing χ to be sufficiently large so that the stability of local updates and the accuracy are satisfactory, i.e., so that they are not significantly improved by increasing χ. For local updates (Sec. IV B 1), we optimize the environment tensors by carrying out sweeps consisting of two CBE sweeps followed by one one-site sweep. For gradient updates (Sec. IV B 2), the virtual basis of the D states is fixed; we therefore choose to keep the basis of the χ states fixed as well and perform only one one-site sweep to optimize the environment.

To initiate the simulations, we use product states made up of patterns of different local states defined either on a site or on a bond. For U(1)-spin-symmetric simulations, we make up the initial product states out of the states $|\uparrow\rangle = |s_z = \tfrac{1}{2}, c_z = 0\rangle$, $|\downarrow\rangle = |s_z = -\tfrac{1}{2}, c_z = 0\rangle$, and $|0\rangle = |s_z = 0, c_z = -\tfrac{1}{2}\rangle$ (the empty local state). For SU(2)-spin-symmetric simulations, we make up the initial product states out of the local states $|0\rangle = |s = 0, c_z = -\tfrac{1}{2}\rangle$ and $|\uparrow\downarrow\rangle = |s = 0, c_z = \tfrac{1}{2}\rangle$ and also out of the bond-singlet state

$$|\text{BS}_{\langle i,j \rangle}\rangle = -\left|s = \tfrac{1}{2}, c_z = 0\right\rangle_i \otimes \left|s = \tfrac{1}{2}, c_z = 0\right\rangle_j, \qquad (44)$$

where $\langle i, j \rangle$ denotes the pair of nearest-neighbor sites i and j. As we will discuss in the following, we find that the convergence behavior of fPEPS does depend on the choice of initial state. We will therefore describe the particular initial product state used for each simulation as well as the effect of the choice of initial state on the convergence. As a reference, we depict product states used to initialize the 3×3 lattice with $\langle n \rangle = 8/9$ for SU(2) as well as for U(1) symmetry in Fig. 32.

In order to investigate the fundamental convergence of the fPEPS algorithm, we start with the minimal case for which all aspects of the algorithm come into play; this is the 3×3 lattice. Furthermore, we first take $U = 0$ as this does not engender trivial convergence of fPEPS so that we can test the convergence behavior and compare the computational cost for a case that is very easy to calculate exactly with almost any numerical method. Note that $U = 0$ is also a nontrivial case for a real-space DMRG calculation.

We first carry out a simulation that only uses local optimization (i.e., no gradient optimization step) in order to explore the efficacy of the local optimization scheme. In
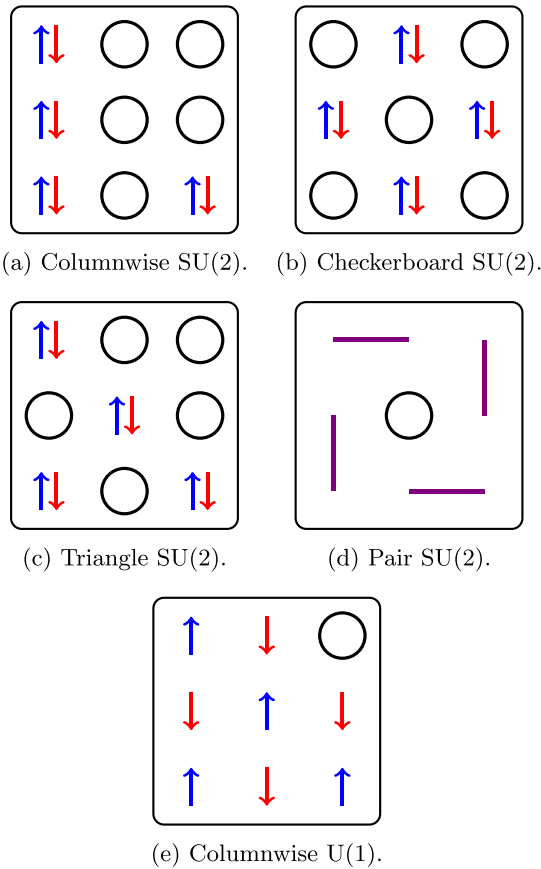
(a) Columnwise SU(2).   (b) Checkerboard SU(2).

(c) Triangle SU(2).   (d) Pair SU(2).

(e) Columnwise U(1).

FIG. 32. Depiction of initial product states for the 3×3 lattice, where states (a), (b), (c), and (d) are used with SU(2) symmetry and state (e) with U(1) symmetry. Here a black circle depicts the empty local state $|0\rangle$, an isolated blue uparrow the local state $|\uparrow\rangle$, an isolated red downarrow the local state $|\downarrow\rangle$, a pair of arrows the doubly occupied local state $|\uparrow\downarrow\rangle$, and a thick purple line the state $|\text{BS}_{\langle i,j\rangle}\rangle$, a bond singlet between nearest-neighbor sites i and j.

FIG. 33. Relative error in the ground-state energy of the Hubbard model on a 3×3 lattice with open boundary conditions for $U = 0$, $S = 0$, and $N = 8$ (i.e., $\langle n \rangle = 8/9$), calculated with fPEPS with SU(2) symmetry using local optimizations only $E_0^{(\text{lu})}$, with respect to the exact ground-state energy E_0 plotted as a function of the sequence of full local update sweeps. For both indicated bond dimensions D, the environment dimension $\chi = 500$.

Fig. 33, we display the convergence of the ground-state energy for 100 full sweeps of local optimization two different bond dimensions, $D = 7$ and $D = 8$, utilizing SU(2) symmetry. Here $N = 8$ electrons, so that $\langle n \rangle = 8/9 \approx 0.89$, a value as close to half filling as is possible for a 3×3 lattice. As the initial state, we take the columnwise SU(2) product state, Fig. 32(a); we will discuss the choice of the initial state in more detail in the context of the combined local and gradient (supersweep) algorithm below. As can be seen, for both values of D, there is systematic convergence to the ground-state energy to a relative accuracy of 10^{-6} as the number of local update full sweeps is increased for both values of D. Furthermore, at least up to approximately 60 update sweeps, there is a systematic improvement of the energy with D. (A possible cause for the poorer convergence of the $D = 8$ calculation relative to the $D = 7$ calculation above 60 sweeps could be that the environment dimension, $\chi = 500$, is too small for the higher bond dimension, $D = 8$.) Particularly notable, however, is how slow the convergence is, especially with the number of sweeps, but also with D. As described in Sec. IV B 1, each full sweep encompasses 18 bond optimizations, where each includes an iterative diagonalization. Thus,
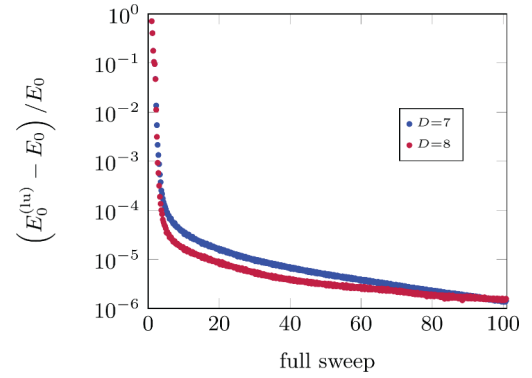
carrying out 100 sweeps is computationally quite expensive—here ca. 8 days of wall time for the $D = 7$ and 58 days for the $D = 8$ runs exclusively using all 16 cores in parallel on a 2.5 GHz Intel Xeon 4215 processor-based compute node. (We will use this compute-node configuration, utilized exclusively and in parallel, as a measure for computational cost for all calculations described in this paper; we will term this measure our "reference compute node".) In contrast, the ground-state energy for this system can be obtained numerically to essentially arbitrary accuracy with exact diagonalization or DMRG in under a second of computer time.

Nevertheless, we emphasize that Fig. 33 demonstrates that systematic variational convergence does occur in fPEPS with local optimization, a result that we have found to be quite hard to achieve. In particular, we have found that if the approximate contraction is not carried out sufficiently carefully (e.g., with a value of χ that is too small), the variationality of the optimization procedure is disturbed, and the variational determination of the energy through iterative diagonalization of the effective Hamiltonian (with a given norm environment), see Fig. 24, becomes unstable.

We now address the question of whether the speed and/or the degree of convergence can be improved by the addition of gradient optimization. Note that the gradient updates described in Sec. IV B 2 optimize all local PEPS tensors A_i simultaneously rather than only two at once for local bond updates and one at once for single-site update (see Sec. IV B 1), but, like single-site updates, cannot change the structure of the basis. Thus, we have found that it is favorable to carry out a mixture of local and gradient optimization, i.e., combinations of the two types of sweeps, which we denote "supersweeps". In particular, we perform successive supersweeps, each consisting of three full local update sweeps followed by 100 gradient updates, repeating until either a saturation of the convergence or a reasonable limit of computer time has been reached. We have empirically found this supersweep configuration to be optimal and use it for all subsequent calculations described below.
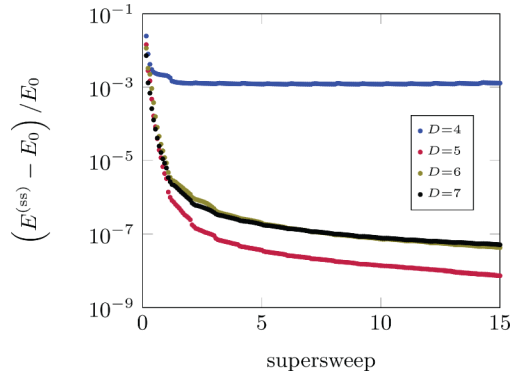
FIG. 34. Relative error in the ground-state energy calculated with fPEPS with SU(2) symmetry using supersweeps consisting of 3 full local update sweeps followed by 100 gradient updates, $E_0^{(ss)}$, with respect to the exact ground-state energy E_0 on the 3×3 Hubbard model with open boundary conditions, $U = 0$, $S = 0$, and $N = 8$ for the indicated values of bond dimension D and $\chi = 500$. The lowest energy for every local update sweep and the energy of every tenth gradient update is plotted.

The results for the relative error in the ground-state energy are plotted in Fig. 34 for values of the bond dimension ranging from $D = 4$ to $D = 7$. As can be seen, there is a regular, systematic convergence with supersweep steps for all values of D. The $D = 4$ curve saturates after approximately two supersweeps (26 update steps), while the curves for higher values of D continue to go down up to the maximum of 15 supersweeps displayed. Going from $D = 4$ to $D = 5$, there is a very large jump in convergence; $D = 4$ saturates at a relative accuracy slightly above 10^{-3}, while the highest relative accuracy of $D = 5$ is below 10^{-8}. Further increasing D, as can be seen from the $D = 6$ and 7 curves, actually reduces the accuracy, with both curves lying approximately on top of one another. Thus, for a $U = 0$ on a 3×3 lattice, there seems to be a breakdown of systematic convergence with D above $D = 5$.

In order to understand the origin of this breakdown, we have examined the state structure of the bond truncation at which the divergence occurs in detail for $D = 5$ and $D = 6$. The D states in the retained set are distributed over various s, c_z quantum numbers, with one or, at most, two states per quantum number. These sets of quantum numbers differ, with a set being retained for $D = 5$ that is not retained for $D = 6$. In view of this, it is our hypothesis that the granularity of the state selection is very sensitive to which states are selected at critical steps; this selection seems to be particularly fortuitous for $D = 5$ in the $U = 0$ case treated here. An examination of the local spin and hole densities, displayed in Fig. 35, shows that the single hole is almost exclusively distributed over the four corner sites of the 3×3 lattice. (The spin and charge densities displayed in Fig. 35 are calculated using fPEPS so that the SU(2) symmetry is explicitly preserved; for $D = 5$, they are essentially numerically exact on the scale of the plot.) This unevenness of this hole distribution is an artifact of the lack of interaction, i.e., that $U = 0$; the distribution of the exact values of the hole densities between the different corner sites is also quite unstable; calculations of lesser accuracy or without SU(2) symmetry result in uneven distributions. Such
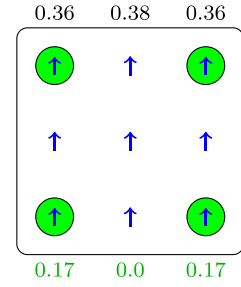

FIG. 35. Local spin density $\langle \mathbf{S}_i^2 \rangle$ (size of blue arrows) and local hole density $1 - \langle n_i \rangle$ (diameter of green-shaded circles) on a 3×3 lattice for $U = 0$, $S = 0$, and $N = 8$, calculated using combined local and gradient updates and SU(2) symmetry with $D = 5$ and $\chi = 500$. The black numbers on the top edge are the averages of the spin densities on the column of sites below, and the green numbers on the bottom edge are the averages of the hole densities on the column of sites above.

an uneven hole distribution that is so sensitive to fine details of the variational state is consistent with the sensitivity to granular state selection described above.

We also note that the convergence behavior is significantly dependent on the initial product state taken. Here we use the checkerboard SU(2) state, Fig. 32(b), which yields the most accurate energies for the $U = 0$, 3×3 system for all D values except $D = 5$. For $D = 5$, using a columnwise SU(2) initial state, Fig. 32(a), yields an energy curve similar to the $D = 5$ curve in Fig. 34, but with a ground-state energy that is approximately a factor of two more accurate (not shown). For all other D values, however, the energies obtained with a columnwise initial state are significantly less accurate than those in Fig. 34. We therefore also ascribe the anomalously accurate energy for the $D = 5$ case with a columnwise initial state to a fortuitous bond truncation similar to that of the checkerboard initial state, but one that happens to be slightly more fortuitous.

We now turn on the interaction to $U = 8$, treating the same system size, 3×3, and band filling, $\langle n \rangle = 8/9$. For this system, we have found that using local optimization alone does not lead to systematic convergence. We therefore only use the supersweep scheme here. However, we do carry out and compare calculations using both the U(1)-symmetric and the SU(2)-symmetric fPEPS states. For the U(1) calculations, we initialize all runs with the columnwise U(1) product state, Fig. 32(e), while for the SU(2) calculations, we use the triangle SU(2) product state, Fig. 32(c). Empirically, we have found these initial states to give the best overall convergence for each symmetry.

We display the convergence of the relative error in the ground-state energy as the algorithm progresses in Fig. 36. For both U(1) and SU(2) symmetry, there is initial rapid convergence for the first few supersweeps, with runs with larger values of D converging to a higher accuracy in that the rapid convergence continues for a larger number of supersweeps. For both symmetries, small wave-like fluctuations are visible at the end of each supersweep, especially for larger values of D. These are the points at which the optimization switches from gradient updates to local updates. At these points, the slope of the convergence of the energy becomes flatter or, in
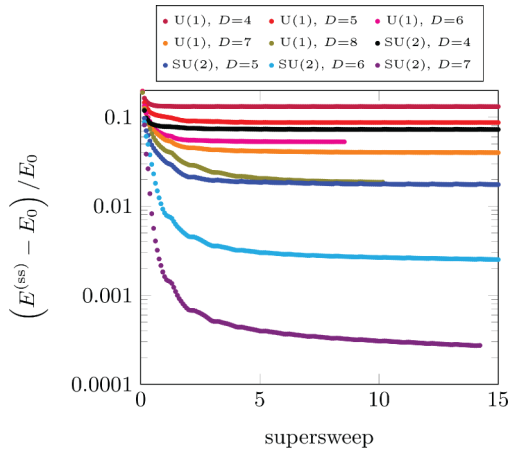
FIG. 36. Relative error in the ground-state energy of the Hubbard model on a 3×3 lattice with open boundary conditions, $U=8$, $S=0$, and $N=8$ ($\langle n\rangle = 8/9$), calculated with fPEPS with U(1) and SU(2) symmetry, as indicated, using supersweeps consisting of three full local optimization sweeps followed by 100 gradient optimizations $E_0^{(\rm ss)}$ with respect to the exact ground-state energy E_0 for $\chi = 500$ and various D. The lowest energy for every local update sweep and the energy of every tenth gradient update is plotted.

FIG. 37. Local spin density $\langle \mathbf{S}_i^2\rangle$ and local hole density $1 - \langle n_i\rangle$ on a 3×3 lattice for $U=8$, $S=0$, and $N=8$, calculated using the DMRG with $D=4000$ states kept and plotted using the same scheme as in Fig. 35.

some cases, particularly where a plateau in the convergence has been reached, actually becomes larger. However, when gradient sweeps begin again, after three full local update sweeps, corresponding to three points, the slope of the convergence becomes steeper again. The behavior demonstrates that neither type of optimization alone displays systematic convergence. The local updates do lead to rapid convergence at the beginning as the basis is adapted from a product state to a more suitable basis for the many-body ground state, but the convergence then slows down significantly. Once a new basis is formed, gradient updates initially lead to a rapid convergence within that basis, but the convergence then saturates, as the basis cannot be changed within the gradient updates. Subsequent local updates do change the local basis, but do not necessarily improve the energy; however, following gradient updates can then continue to improve the energy until saturation is reached once again. Note that this behavior can also be seen for the $U=0$ case in Fig. 34, albeit at a slightly smaller scale so that it is not as visible.

For the U(1) calculations, the relative error reaches successively lower plateaus for all values of D with the largest value, $D=8$, reaching a value of approximately 0.02 after five supersweeps. The SU(2) calculations also have an initial rapid convergence for all D values that crosses over to either a plateau for $D=4$ and $D=5$ or a more gradual decline for $D=6$ and $D=7$ on a scale of approximately five supersweeps. The accuracy improves systematically and rapidly with D. The irregular convergence with D seen for the $U=0$ calculations with SU(2) symmetry in Fig. 34, where accuracy in the energy is much higher for $D=5$ than for $D=6$ and $D=7$, is not present here. We observe that the distribution of charge and spin for the $U=8$ case, displayed in Fig. 37 (here calculated numerically exactly using the DMRG), is much more uniform than that for the $U=0$ case, Fig. 35. Evidently, turning on the interaction leads to a much more
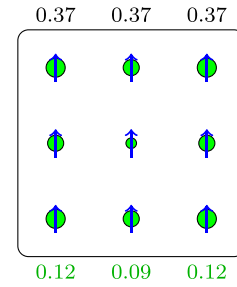
uniform distribution and, presumably, reduces the problems due to granularity in state selection within the bond truncation. Note, however, that the relative error in the ground-state energy is, in fact, much smaller in best case for $U=0$, $D=5$, than in the most accurate $U=8$ calculation, which has the highest bond dimension, $D=7$.

In all cases, a particular SU(2) calculation is significantly more accurate than the U(1) calculation with the same D. This significant improvement in going from U(1) to SU(2) symmetry is expected, as one SU(2) state comprises an entire spin multiplet, as discussed in Sec. III. Thus, a given SU(2) bond dimension D effectively encompasses a region of the total Hilbert space that could only be encompassed with a U(1)-symmetric fPEPS with a significantly higher bond dimension. We also note that SU(2)-symmetric calculations explicitly preserve the fundamental spin-rotation symmetry present in the system, which is not the case for U(1) calculations.

The minimum relative error reached, for the $D=7$, SU(2) calculation, is just under 3×10^{-4}. This calculation took approximately nine days of wall time on our reference compute node. Thus, the fPEPS method is certainly not competitive either in accuracy or in computational efficiency with other numerical methods that are numerically exact on this small system, such as exact diagonalization and the DMRG.

We now turn to larger system sizes. Based on the results for the 3×3 system discussed above, we expect the computational effort needed to obtain reasonable convergence, i.e., at least to the level of Fig. 36, to be unreachably large for larger lattice sizes. Nevertheless, we have carried out runs that are as long as practicably possible in order to gain knowledge about the convergence behavior.

We first consider the 4×4 lattice, whose ground state is still accessible using exact diagonalization. Our primary purpose here (and for the 6×6 lattice treated below) is to test the convergence behavior. Thus, we try to minimize detailed physical effects due to charge inhomogeneity and treat the half-filled system, i.e., take $\langle n\rangle = 1$, that is, $N=16$ electrons on the 4×4 lattice. This choice also simplifies the choice of the initial state: for U(1) symmetry, we start the simulation with a Néel state, i.e., a product state with antiferromagnetically alternating local spins, a pattern corresponding to the spin pattern in Fig. 32(e) (i.e., with no hole). For the SU(2) calculations, we take a columnwise SU(2) state, i.e., the half-filled analog of Fig. 32(a) on a 4×4 lattice, that is, two columns of
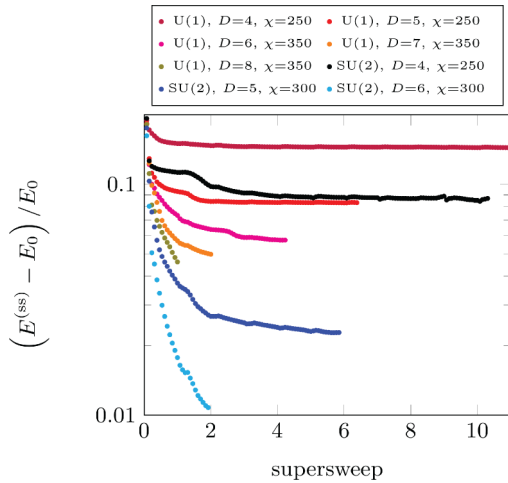
FIG. 38. Relative error in the ground-state energy of the Hubbard model on a 4×4 lattice with open boundary conditions, $U = 8$, $S = 0$, and $N = 16$ (half filling), calculated with fPEPS with U(1) and SU(2) symmetry, as indicated, using supersweeps consisting of three full local optimization sweeps followed by 100 gradient optimizations $E_0^{(ss)}$ with respect to the exact ground-state energy E_0 for various D with the indicated χ. The lowest energy for every local update sweep and the energy of every tenth gradient update is plotted.

FIG. 39. Relative error in the ground-state energy of the Hubbard model on a 6×6 lattice with open boundary conditions, $U = 8$, $S = 0$, and $N = 36$ (half filling), calculated with fPEPS with U(1) and SU(2) symmetry, as indicated, using supersweeps consisting of three full local optimization sweeps followed by 100 gradient optimizations $E_0^{(ss)}$ with respect to the ground-state energy calculated using the DMRG E_0, for various D with the indicated χ. The lowest energy for every local update sweep and the energy of every tenth gradient update is plotted.

doubly occupied local states interspersed with two columns of unoccupied states. We have ascertained that choosing a physically reasonable alternate initial product state, in particular, one in which the lattice is covered by vertically oriented bond-singlet pairs, cf. Fig. 32(d), does not significantly effect the convergence behavior. In this sense, fPEPS for this larger, half-filled system seems to be significantly less sensitive to choice of initial states than the 3×3 systems with one hole discussed above.

Figure 38 depicts the the relative error in the ground-state energy on a half-filled 4×4 lattice at $U = 8$ as a function of the optimization step for runs with both U(1) and SU(2) symmetry and various bond dimensions ranging from $D = 4$ to $D = 8$. Consistent with the behavior for the 3×3, $U = 8$ system, there is systematic convergence both with increasing supersweep step and with increasing D. As before, utilizing SU(2) rather than U(1) symmetry for a given fixed D also leads to a significant improvement in the relative accuracy for the 4×4 system. The best relative accuracy in energy of about 0.01 is obtained from the SU(2) simulation with $D = 6$ states, which took approximately four days of wall time on our reference compute node; this is the largest bond dimension that we could practically retain for this system. Taking the 3×3, $U = 8$ calculations as a guide, at least 15 supersweeps would likely be required to achieve reasonable convergence for $D = 6$. This would correspond to a wall time of approximately one month on our reference compute node.

Figure 39 displays corresponding results for the accuracy of the ground-state energy on a 6×6 lattice relative to DMRG calculations. (Note that the DMRG energy is no longer numerically exact on a 6×6 lattice, but should be accurate to at least 2 or 3 significant digits more than the best fPEPS calculation.) The initial states are scaled-up versions of those used for the half-filled 4×4 calculations: columnwise states
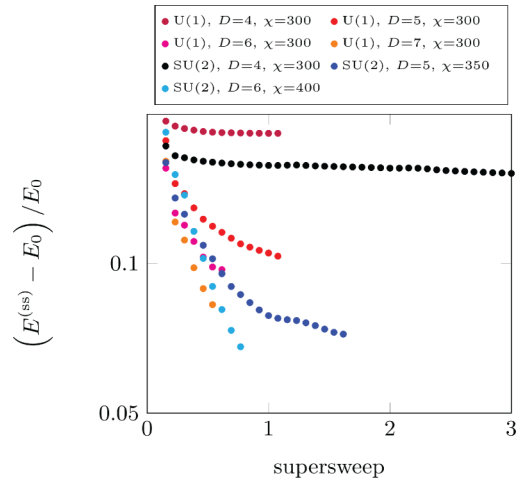
with alternating columns of double- and zero-occupied sites for SU(2) symmetry and a Néel state for U(1) symmetry. The overall features of the energy convergence found in the calculations for $U = 8$ on the 3×3 and 4×4 lattices are retained for this larger system size: there is systematic convergence in supersweep step and in the energy convergence curves as D is increased. For a given bond dimension D, the SU(2) calculations are substantially more accurate than the U(1) calculations, albeit at a higher computational cost. Due to the cost of the calculations, none of the runs could practicably be carried out for more than two supersweeps except for the computationally least expensive $D = 4$, U(1) calculation. The best accuracy of the energy relative to DMRG achieved in this scope, for the $D = 6$, SU(2) calculations, is approximately 7%, almost an order of magnitude less accurate than the best accuracy obtained for the 4×4 lattice. Note that, as for the 4×4 system, we vary χ from run to run in order to balance stability and accuracy with computational cost. In particular, the $D = 5$ and $D = 6$, SU(2)-symmetric runs require higher values of χ to ensure stability than runs with the corresponding D for the 4×4 lattice. As on the 4×4 lattice, the SU(2), $D = 6$ case is far from converged in the number of supersweeps. The wall time needed for the calculation shown, which comprises less than one full supersweep, is approximately 14 days, so that reasonable convergence in the number of supersweeps would take at least several months if not years of wall time on our reference compute node. While it would be useful to be able to extract systematics in the scaling of accuracy and computational cost with system size D and χ for half-filled systems, it is not practical to do this here, as we have not been able to achieve sufficient convergence with reasonable computational resources for the two smallest system sizes;
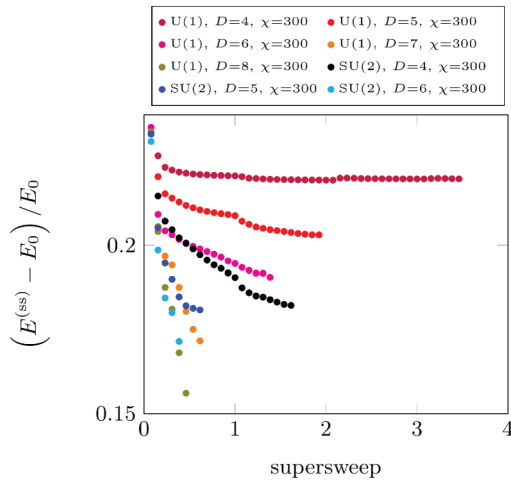
FIG. 40. Relative error in the ground-state energy for the Hubbard model on an 8×8 lattice for $U=8$, $S=0$, and $N=56$ ($\langle n\rangle = 0.875$), calculated with fPEPS with U(1) and SU(2) symmetry, as indicated, using supersweeps consisting of three full local optimization sweeps followed by 100 gradient optimizations $E_0^{(\mathrm{ss})}$ with respect to the ground-state energy from a DMRG calculation (see text) E_0. For every local update sweep, the lowest energy is plotted, and the energy of every tenth gradient update is plotted.

treating larger systems is clearly outside of reasonable scope for the computational resources available to us.

However, in order to explore to what extent the fPEPS method can capture the essential physics of the doped two-dimensional Hubbard model qualitatively, we nevertheless present here simulation results for the 8×8 lattice at 1/8-doping, i.e., with $N=56$ electrons so that $\langle n\rangle = 0.875$. The simulations were initialized with scaled-up versions of the pair SU(2) state, Fig. 32(d), for the SU(2) simulations and with scaled-up versions of columnwise U(1) state for the U(1) simulations. We have found that the choice of initial state does have a substantial effect on the initial convergence of the SU(2) simulations especially; for example, the accuracy is substantially lower if a columnwise SU(2) state, i.e., a scaled-up version of Fig. 32(a), is used for the SU(2) simulations. Figure 40 displays the error in the ground-state energy relative to DMRG calculations. Here we expect the DMRG calculations to have substantially lower relative accuracy than then 6×6 calculations, so that we estimate that the relative error in the ground-state energy obtained using the DMRG could be up to the order of 1%. However, as can be seen in Fig. 40, the deviation of the ground-state energy calculated with fPEPS from that calculated using the DMRG does not go below 10%. In this sense, it is not necessary to have more accurate variational estimates of the ground-state energy for comparison. Surprisingly, we find that a slightly lower χ, as compared to the 6×6, half-filled system, is sufficient for stable SU(2) simulations, despite the fact that the system size is larger. Thus, doping seems to lead to greater stability in the variational optimization for given D and χ. The origins of this behavior are unclear: while the Hilbert space of the doped system is smaller than the half-filled system, the 1/8-doped 8×8 system should nevertheless have a larger Hilbert space than the half-filled 6×6 system. It could be that the breaking

of translational invariance in both the charge and spin sectors due to stripe structures (see below) helps increase the stability of the calculations.

The smaller environment dimensions χ required allow us to carry out calculations with comparable values of D [including even a higher value, $D=8$, in the U(1) calculations] and a comparable number of supersweep steps as for the half-filled 6×6 system. The convergence behavior is qualitatively similar to that of the $U=8$ calculations on all other system sizes in that the convergence systematically improves as a function of supersweep step, of D, and in going from U(1) to SU(2) symmetry. The wavelike variations at the boundaries between supersweeps is also present. Note, however, that the lowest ground-state energy obtained (actually for the $D=8$ U(1) calculation) is about 10% above the DMRG energy. The wall time required for the $D=8$, U(1) calculation, which consists of approximately half of a supersweep, was a little more than 21 days on our reference compute node. For the $D=6$, SU(2) calculation, the wall time was almost 29 days.

In order to gain insight into the convergence behavior and understand to what extent the fPEPS calculations can capture the physical behavior of the doped two-dimensional Hubbard model qualitatively, we now examine the local spin and hole densities for both U(1) and SU(2) calculations. We note that doing this on the 8×8 rather than on smaller lattices is appropriate because the ground state is expected to have a stripe structure with wavelength $\lambda = 8$ stripes [14]. Such a stripe structure would be frustrated on smaller lattices, leading to complicated patterns in the local spin and hole densities and, potentially, to instabilities in the energetically most favorable configurations.

We depict the local spin density $\langle S_i^z\rangle = \frac{1}{2}(n_{i,\uparrow} - n_{i,\downarrow})$ and local hole density $1-\langle n_i\rangle$ graphically on the 8×8 lattice for the most accurate U(1)-symmetric calculation, the one with $D=8$ and $\chi=300$, in Fig. 41. [Since the U(1) calculation breaks spin-rotation symmetry, the magnetic structure can best be seen using maps of $\langle S_i^z\rangle$.] The stripe structures found in Ref. [14] using a variety of independent, state-of-the-art numerical methods is well reproduced: There is a vertical site-centered hole stripe separating antiferromagnetic regions with an antiphase boundary between the regions that has a total wavelength of $\lambda = 8$. Note that this state is significantly evolved from the initial state, which is a product state with a Néel configuration in the first seven columns and a vertical line of holes in the eighth column. If one examines the corresponding spin and hole densities as the convergence progresses, i.e., as more supersweep steps are carried out, one can see that the initial configuration evolves step-by-step to the configuration depicted in Fig. 41 in that the line of holes on the eighth column are moved to the left and spread out and that the antiferromagnetic spin structure with opposite phase forms to the right of the stripe. Stripe structures intermediate between the initial product state and those of Fig. 41 are also found in U(1) calculations with smaller bond dimension D.

We also examine the spin and hole densities for the most accurate SU(2)-symmetric calculations, for $D=6$ and $\chi=300$, in Fig. 42. Since $\langle S_i^z\rangle = 0$ identically on all sites i for SU(2)-symmetric calculations, we instead depict the local spin moment $\langle \mathbf{S}_i^2\rangle$, which is large in magnetically ordered
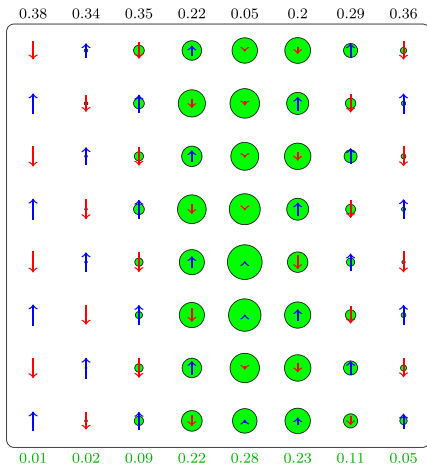
FIG. 41. Local z component of the spin $\langle S_i^z \rangle = \frac{1}{2}(n_{i,\uparrow} - n_{i,\downarrow})$ (size, color, and direction of arrows) and local hole density $1 - \langle n_i \rangle$ (diameter of green-shaded circles) on an 8×8 lattice with open boundary conditions calculated with U(1) symmetry, bond dimension $D = 8$, and $\chi = 300$. Here $U = 8$, $S_z = 0$, and $N = 56$ so that $\langle n \rangle = 0.875$. The black numbers are the average $\langle S_i^z \rangle$ for the column of sites below, and the green numbers on the bottom edge are the average hole densities for the column of sites above.

regions, but cannot depict the antiferromagnetic structure explicitly. (This could be done by calculating spin-spin correlations between pairs of sites.) The initial product state for the SU(2) calculation is a pair-bond state of the type depicted in Fig. 32(d) for the 3×3 lattice, but with vertical bond-singlet states covering the first seven columns and a vertical line of zero-occupied local states in the eighth column. As can be seen in Fig. 42, this state evolves towards a stripe configuration roughly consistent with that of the U(1) calculation,
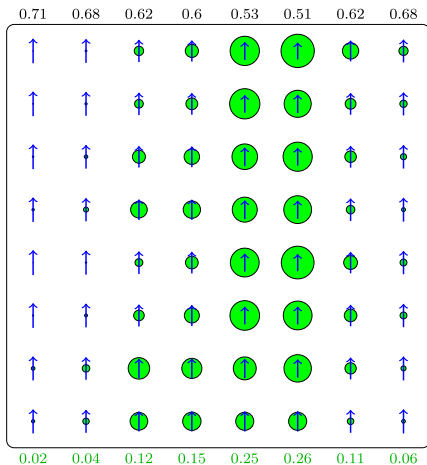

FIG. 42. Local spin density $\langle \mathbf{S}_i^2 \rangle$ (size of blue arrows) and local hole density $1 - \langle n_i \rangle$ (diameter of green-shaded circles) for the Hubbard model on an 8×8 lattice with open boundary conditions and $U = 8$, $S = 0$, and $N = 56$ so that $\langle n \rangle = 0.875$, calculated with SU(2) symmetry, bond dimension $D = 6$, and $\chi = 300$. The black numbers are the average spin density for the column of sites below, and the green numbers on the bottom edge are the average hole densities for the column of sites above.

Fig. 41. We note that the distribution of holes is, however, not quite as symmetric in both the horizontal and in the vertical directions as that of the U(1) calculation and that the hole stripe is not as clearly site centered on the fifth column of sites. It is important to realize that this calculation is fairly far from convergence in that only three local update sweeps and 30 gradient updates have been carried out within the first supersweep and that the energy is substantially higher than the lowest energy for the U(1), $D = 8$ calculation in Fig. 40. Thus, we ascribe these discrepancies to the fact that the SU(2) calculation is, relatively speaking, more poorly converged than the U(1) calculation.

VI. SUMMARY AND DISCUSSION

In this paper, we have developed the groundwork to carry out variational optimization of finite projected entangled pair states. We treat the two-dimensional Hubbard model in this paper because we are interested in developing methods for short-range itinerant fermionic models in which a high degree of symmetry is present. The single-band model with nearest-neighbor hopping displays interesting and complicated physical behavior, which is partially, but not fully, understood, and is accessible to a set of other numerical methods, albeit at a level that pushes the limits of these methods. In this sense, it is suited to be a stringent and demanding test bed for our methods. However, we emphasize that the fPEPS method is applicable to wide set of short-ranged, two-dimensional quantum lattice models: in particular, to quantum-spin-based models such as Heisenberg and t-J models, to spinless-fermion models, as well as to extended versions of the Hubbard model that include features such as longer-range hopping or interaction and multiple bands.

Our method is based on the following essential building blocks: (i) a general framework to create and handle projected entangled pair operators for arbitrary local Hamiltonians (Sec. II B), (ii) a generic scheme to incorporate SU(2)$_{\text{spin}}$ symmetry (Sec. III) into PEPS states at the level of tensor representations, (iii) a procedure, which optimizes the environment within PEPS-based contractions using a two-block configuration (Sec. IV A), and (iv) the utilization of a generalized version of the controlled bond expansion [58] to circumvent contracting large tensors. We have integrated these technical contributions with methods and representations from the extensive existing body of knowledge about tensor networks in order to formulate a comprehensive, PEPO-based framework for calculating ground-state properties of two-dimensional quantum lattice systems.

Given this framework, which enables us to represent variational states as finite PEPS and operators on these states as PEPOs, the crucial component to build an algorithm to approximate the ground state is a method to optimize the fPEPS tensor network for a given Hamiltonian, i.e., given a PEPO. Here we have formulated two optimization methods: a local-update optimization procedure that optimized two adjacent local tensors within the fPEPS via iterative diagonalization and a gradient-update optimization scheme that optimizes the elements of all local tensors simultaneously. We find that the methods are complementary in that local optimization can change the structure of the basis (i.e., what states with which

quantum numbers are included) on a particular bond between local tensors in addition to optimizing the elements of these two tensors, whereas gradient optimization carries out its simultaneous optimization of all local tensor elements within a fixed basis of all tensor indices.

Test simulations, presented in Sec. V, show that our implementation of the fPEPS algorithm does, in fact, achieve systematic variational improvement in the number of optimization steps if combined local and gradient optimizations are carried out (what we call the "supersweep" procedure), as long as care is taken that the environment dimension χ is chosen to be high enough for all systems and that a suitable initial state is chosen. For small values of the bond dimension D, the convergence saturates after a certain number of optimization steps; for the larger values of D, such a saturation does not occur on the scale of computer time accessible to our calculations. For all systems with nonzero Hubbard interaction U, increasing the bond dimension D for calculations with either U(1) or SU(2) symmetry systematically lowers the variational approximation to the ground-state energy. Note that this is not so for the 3×3, $U=0$ test case, for which the best numerical accuracy is achieved for bond dimension $D=5$, and further increasing D actually increases the approximate ground-state energy. A detailed analysis suggests, however, that this behavior can likely be ascribed to particularities of the $U=0$ state with one hole, namely, a very inhomogeneous distribution of the single hole over the corner sites and a near-degeneracy to variations of the hole density on these sites. For all systems, calculations with SU(2) symmetry with a given bond dimension D are substantially more accurate than U(1)-symmetric calculations with the same D. This is as expected because a single SU(2) state, in general corresponds to a multiplet of states in the U(1) representation. Note, however, that an SU(2) calculation is generally substantially more expensive computationally, both in computer time and in memory, than a U(1) calculation with the same D. Nevertheless, for all systems except for 8×8, $U=8$, we have achieved the most accurate results using SU(2) symmetry.

Despite these generally satisfying aspects of the convergence behavior, the results show that fPEPS is not competitive with other numerical methods, in particular, with MPS-based method such as the DMRG, in calculating the ground state of the two-dimensional Hubbard model and other two-dimensional quantum lattice models. The main obstacle seems to be the poor convergence of the optimization process. Except for $U=0$, local updates alone do not seem to be able to reach the lowest possible energy for a given bond dimension D. Adding gradient updates within the supersweep procedure overcomes this barrier, but the variational ground-state energy for a given D only approaches the exact ground state energy E_0 after an unreasonably large number of optimization steps, as is exemplified for the 3×3, $U=8$ test system in Fig. 36; the runs with larger values of D do not converge in optimization steps at all for reachable computer time, even for this small test system. For larger systems, we do observe that increasing D leads to lower energies, but, due to the computational cost of carrying out these calculations, it is hard to estimate at what point and at what level of accuracy of the energy would converge in these runs. The best relative error in the ground-state energy reached is, in absolute terms,

quite bad for a variational calculation: ca. 1% for the 4×4 and ca. 7% for the 6×6, $U=8$, half-filled systems, and ca. 10% for the 8×8, $U=8$, 1/8-doped systems. Nevertheless, the results do display some aspects that are encouraging: for the lowest values of D, there does seem to be convergence in optimization steps, even for the 8×8 system [at least with U(1) symmetry]. The behavior of the spin and hole densities of the doped Hubbard model on the 8×8 lattice, Figs. 41 and 42, is in qualitative agreement with the expected stripe configuration known from calculations with other methods, both for U(1) symmetry, and, to a lesser extent because of the rather incomplete convergence, for the SU(2) symmetry, despite the relatively inaccurate ground-state energies.

Thus, we are lead to the question of whether it is possible to substantially improve the convergence behavior of fPEPS. This question can be divided into two aspects: First, how well does an fPEPS approximate a particular many-body ground state on a fundamental level? Second, can existing optimization procedures be improved or alternative optimization procedures be developed so that the optimal fPEPS for a given bond dimension D can efficiently be found? As to the first question, our results have both encouraging and cautionary aspects. Encouraging is that increasing D or going from U(1) to SU(2) symmetry does seem to systematically improve the fPEPS approximation to the ground state. How this improvement takes place and what D is required to obtain a given accuracy for given Hamiltonian parameters, particle number, and system size could only be explored to a limited extent due to the poor convergence for all but the smallest systems. Cautionary is the rather sensitive dependence of not only the rate of convergence, but also the level of apparent convergence after many optimization steps, on the initial state.

As to the second question, the fact that both local optimization and gradient optimization have features that work in a complementary fashion indicates that neither method alone is sufficient to carry out an efficient optimization. Successive local optimizations of adjacent sites do not seem to lead towards a minimum of the entire fPEPS, whereas gradient-based optimization is, by construction, unable to modify the bases on the local tensor bonds to form an fPEPS that efficiently approximates the many-body ground state. Combining the methods partially overcomes these problems, but the synergy between the two methods is limited, so that convergence is still very slow. A more effective optimization algorithm would presumably have to combine the basis adaptation of the local updates and the nonlocal optimization of PEPS tensors of the gradient updates in a coherent way.

Note that a large amount of previous work on both fPEPS and iPEPS algorithms uses imaginary time evolution to optimize the PEPSs. Usually, this is done by carrying out a Trotter decomposition and applying terms of the Hamiltonian locally on sites and bonds. We therefore expect imaginary time evolution to have, at best, a convergence along the lines of, but poorer than that of our local optimization scheme. In addition, imaginary time evolution introduces an additional systematic Trotter error.

While the DMRG may still be faster and more accurate than fPEPS for all system sizes treated here, it is limited in scaling two-dimensional systems to the thermodynamic limit

due to the exponential increase of states needed to maintain accuracy as the lattice width is increased, which is due to the entropy area law [35]. PEPSs, on the other hand, are, in principle, capable of efficiently describing states of arbitrarily large two-dimensional lattices as long as they satisfy the entropy area law [37]. Thus, should it be possible to overcome the poor convergence of fPEPS using a better optimization scheme, fPEPS could potentially be competitive with MPS-based methods as two-dimensional lattices are scaled to larger size. In principle, one would like to compare the results of fPEPS simulations to those of iPEPS simulations [79]. Since an iPEPS represents, by construction, a wavefunction in the thermodynamic limit, the crucial step for bringing the fPEPS results onto the same footing would be a finite-size extrapolation of the fPEPS results. Note that such an extrapolation is necessary to compare any numerically exact finite-lattice results, such as those from exact diagonalization or DMRG, to infinite-system results; a comparison of such extrapolated results for the Hubbard model has already been carried out in Ref. [14]. For both iPEPS and fPEPS, extrapolation in bond dimension D and environment dimension χ should also be carried out; how best to do this for imaginary-time-evolved iPEPS results is discussed extensively in Ref. [79].

An additional possible direction for further development is in applying the methods for representing, manipulating, and optimizing PEPS-like states developed in this paper to the iPEPS algorithm. The iPEPS method carries out the approximate contraction of the PEPS tensor network using a corner-transfer matrix rather than the row-wise contraction into an effective MPS used for fPEPS, so that a number of aspects of the methods developed here would have to be adapted to the iPEPS contraction scheme. iPEPS usually works with either a single translationally invariant local PEPS tensor or with a unit cell of such tensors of limited size, used when the physics of the system is expected to break translational invariance. The convergence problems in fPEPS occur only for system sizes of 3×3 or larger because in smaller fPEPS, the environment reduces to what is essentially a pure MPS. It would be interesting to investigate the convergence behavior of local and gradient optimization methods when applied to iPEPS when a unit cell of 3×3 independent local tensors or more is used.

[1] J. Hubbard, Electron correlations in narrow energy bands, Proc. R. Soc. A **276**, 238 (1963).

[2] N. F. Mott and R. W. Gurney, *Electronic Processes in Ionic Crystals* (Clarendon Press, Oxford, 1948).

[3] N. F. Mott, Metal-insulator transition, Rev. Mod. Phys. **40**, 677 (1968).

[4] J. G. Bednorz and K. A. Müller, Possible high T_c superconductivity in the Ba-La-Cu-O system, Z. Phys. B **64**, 189 (1986).

[5] D. J. Scalapino, E. Loh, and J. E. Hirsch, d-wave pairing near a spin-density-wave instability, Phys. Rev. B **34**, 8190(R) (1986).

[6] D. J. Scalapino, E. Loh, and J. E. Hirsch, Fermi-surface instabilities and superconducting d-wave pairing, Phys. Rev. B **35**, 6694 (1987).

[7] D. J. Scalapino, The case for $d_{x^2-y^2}$ pairing in the cuprate superconductors, Phys. Rep. **250**, 329 (1995).

[8] P. Monthoux, A. V. Balatsky, and D. Pines, Toward a Theory of High-Temperature Superconductivity in the Antiferromagnetically Correlated Cuprate Oxides, Phys. Rev. Lett. **67**, 3448 (1991).

[9] S. R. White, D. J. Scalapino, R. L. Sugar, N. E. Bickers, and R. T. Scalettar, Attractive and repulsive pairing interaction vertices for the two-dimensional Hubbard model, Phys. Rev. B **39**, 839 (1989).

[10] J. M. Tranquada, B. J. Sternlieb, J. D. Axe, Y. Nakamura, and S. Uchida, Evidence for stripe correlations of spins and holes in copper oxide superconductors, Nature (London) **375**, 561 (1995).

[11] J. M. Tranquada, J. D. Axe, N. Ichikawa, Y. Nakamura, S. Uchida, and B. Nachumi, Neutron-scattering study of stripe-phase order of holes and spins in La$_{1.48}$Nd$_{0.4}$Sr$_{0.12}$CuO$_4$, Phys. Rev. B **54**, 7489 (1996).

[12] J. M. Tranquada, J. D. Axe, N. Ichikawa, A. R. Moodenbaugh, Y. Nakamura, and S. Uchida, Coexistence of, and Competition between, Superconductivity and Charge-Stripe Order in La$_{1.6-x}$Nd$_{0.4}$Sr$_x$CuO$_4$, Phys. Rev. Lett. **78**, 338 (1997).

[13] S. R. White and D. J. Scalapino, Ground states of the doped four-leg t-J ladder, Phys. Rev. B **55**, R14701 (1997).

[14] B.-X. Zheng, C.-M. Chung, P. Corboz, G. Ehlers, M.-P. Qin, R. M. Noack, H. Shi, S. R. White, S. Zhang, and G. K.-L. Chan, Stripe order in the underdoped region of the two-dimensional Hubbard model, Science **358**, 1155 (2017).

[15] D. P. Arovas, E. Berg, S. A. Kivelson, and S. Raghu, The Hubbard model, Annu. Rev. Condens. Matter Phys. **13**, 239 (2022).

[16] M. Qin, T. Schäfer, S. Andergassen, P. Corboz, and E. Gull, The Hubbard model: A computational perspective, Annu. Rev. Condens. Matter Phys. **13**, 275 (2022).

[17] S. R. White, Density Matrix Formulation for Quantum Renormalization Groups, Phys. Rev. Lett. **69**, 2863 (1992).

[18] S. R. White, Density-matrix algorithms for quantum renormalization groups, Phys. Rev. B **48**, 10345 (1993).

[19] S. Östlund and S. Rommer, Thermodynamic Limit of Density Matrix Renormalization, Phys. Rev. Lett. **75**, 3537 (1995).

[20] R. M. Noack, S. R. White, and D. J. Scalapino, Correlations in a Two-Chain Hubbard Model, Phys. Rev. Lett. **73**, 882 (1994).

[21] S. R. White, R. M. Noack, and D. J. Scalapino, Resonating Valence Bond Theory of Coupled Heisenberg Chains, Phys. Rev. Lett. **73**, 886 (1994).

[22] R. M. Noack, S. R. White, and D. J. Scalapino, The doped two-chain Hubbard model, Europhys. Lett. **30**, 163 (1995).

[23] R. M. Noack, S. R. White, and D. J. Scalapino, The ground state of the two-leg Hubbard ladder a density-matrix renormalization group study, Physica C: Superconductivity **270**, 281 (1996).

[24] S. R. White and D. J. Scalapino, Density Matrix Renormalization Group Study of the Striped Phase in the 2D $t − J$ Model, Phys. Rev. Lett. **80**, 1272 (1998).

[25] R. M. Noack, N. Bulut, D. J. Scalapino, and M. G. Zacher, Enhanced $d_{x^2−y^2}$ pairing correlations in the two-leg Hubbard ladder, Phys. Rev. B **56**, 7162 (1997).

[26] S. R. White and D. J. Scalapino, Energetics of Domain Walls in the 2D $t − J$ Model, Phys. Rev. Lett. **81**, 3227 (1998).

[27] S. R. White and D. J. Scalapino, Competition between stripes and pairing in a $t − t' − J$ model, Phys. Rev. B **60**, R753(R) (1999).

[28] S. Yan, D. A. Huse, and S. R. White, Spin-Liquid Ground State of the $S=$ 1/2 kagome Heisenberg antiferromagnet, Science **332**, 1173 (2011).

[29] E. M. Stoudenmire and S. R. White, Studying two-dimensional systems with the density matrix renormalization group, Annu. Rev. Condens. Matter Phys. **3**, 111 (2012).

[30] S. Depenbrock, I. P. McCulloch, and U. Schollwöck, Nature of the Spin-Liquid Ground State of the $S = 1/2$ Heisenberg Model on the Kagome Lattice, Phys. Rev. Lett. **109**, 067201 (2012).

[31] Simons Collaboration on the Many-Electron Problem, J. P. F. LeBlanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, E. Kozik, X.-W. Liu, A. J. Millis, N. V. Prokof'ev, M. Qin, G. E. Scuseria, H. Shi, B. V. Svistunov, L. F. Tocchio *et al.*, Solutions of the Two-Dimensional Hubbard Model: Benchmarks and Results from a Wide Range of Numerical Algorithms, Phys. Rev. X **5**, 041041 (2015).

[32] G. Ehlers, J. Sólyom, Ö. Legeza, and R. M. Noack, Entanglement structure of the Hubbard model in momentum space, Phys. Rev. B **92**, 235116 (2015).

[33] G. Ehlers, S. R. White, and R. M. Noack, Hybrid-space density matrix renormalization group study of the doped two-dimensional Hubbard model, Phys. Rev. B **95**, 125125 (2017).

[34] G. Ehlers, B. Lenz, S. R. Manmana, and R. M. Noack, Anisotropy crossover in the frustrated Hubbard model on four-chain cylinders, Phys. Rev. B **97**, 035118 (2018).

[35] M. B. Hastings, An area law for one-dimensional quantum systems, J. Stat. Mech. (2007) P08024.

[36] M. M. Wolf, F. Verstraete, M. B. Hastings, and J. I. Cirac, Area Laws in Quantum Systems: Mutual Information and Correlations, Phys. Rev. Lett. **100**, 070502 (2008).

[37] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions, arXiv:cond-mat/0407066.

[38] D. Perez-Garcia, F. Verstraete, J. I. Cirac, and M. M. Wolf, PEPS as unique ground states of local Hamiltonians, arXiv:0707.2260.

[39] N. Schuch, I. Cirac, and D. Pérez-García, PEPS as ground states: Degeneracy and topology, Ann. Phys. **325**, 2153 (2010).

[40] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Criticality, the Area Law, and the Computational Power of Projected Entangled Pair States, Phys. Rev. Lett. **96**, 220601 (2006).

[41] V. Murg, F. Verstraete, and J. I. Cirac, Variational study of hardcore bosons in a two-dimensional optical lattice using projected entangled pair states, Phys. Rev. A **75**, 033605 (2007).

[42] V. Murg, F. Verstraete, and J. I. Cirac, Exploring frustrated spin systems using projected entangled pair states, Phys. Rev. B **79**, 195119 (2009).

[43] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Unifying projected entangled pair state contractions, New J. Phys. **16**, 033014 (2014).

[44] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Algorithms for finite projected entangled pair states, Phys. Rev. B **90**, 064425 (2014).

[45] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, Classical Simulation of Infinite-Size Quantum Lattice Systems in Two Spatial Dimensions, Phys. Rev. Lett. **101**, 250602 (2008).

[46] P. Corboz, R. Orús, B. Bauer, and G. Vidal, Simulation of strongly correlated fermions in two spatial dimensions with fermionic projected entangled-pair states, Phys. Rev. B **81**, 165104 (2010).

[47] P. Corboz, S. R. White, G. Vidal, and M. Troyer, Stripes in the two-dimensional t-J model with infinite projected entangled-pair states, Phys. Rev. B **84**, 041108(R) (2011).

[48] P. Corboz, J. Jordan, and G. Vidal, Simulation of fermionic lattice models in two dimensions with projected entangled-pair states: Next-nearest neighbor Hamiltonians, Phys. Rev. B **82**, 245119 (2010).

[49] P. Corboz, T. M. Rice, and M. Troyer, Competing States in the t-J Model: Uniform d-Wave State versus Stripe State, Phys. Rev. Lett. **113**, 046402 (2014).

[50] L. Vanderstraeten, J. Haegeman, P. Corboz, and F. Verstraete, Gradient methods for variational optimization of projected entangled-pair states, Phys. Rev. B **94**, 155123 (2016).

[51] W.-Y. Liu, S.-J. Dong, Y.-J. Han, G.-C. Guo, and L. He, Gradient optimization of finite projected entangled pair states, Phys. Rev. B **95**, 195154 (2017).

[52] S.-J. Dong, C. Wang, Y. Han, G.-C. Guo, and L. He, Gradient optimization of fermionic projected entangled pair states on directed lattices, Phys. Rev. B **99**, 195153 (2019).

[53] To our knowledge, the PEPO idea was first introduced by Crosswhite in Ref. [62].

[54] P. Corboz and G. Vidal, Fermionic multiscale entanglement renormalization ansatz, Phys. Rev. B **80**, 165129 (2009).

[55] P. Corboz, Variational optimization with infinite projected entangled-pair states, Phys. Rev. B **94**, 035133 (2016).

[56] K. Hyatt and E. M. Stoudenmire, DMRG approach to optimizing two-dimensional tensor networks, arXiv:1908.08833.

[57] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu *et al.*, Is there evidence for exponential quantum advantage in quantum chemistry? arXiv:2208.02199.

[58] A. Gleis, J.-W. Li, and J. von Delft, Controlled bond expansion for DMRG ground state search at single-site costs, arXiv:2207.14712.

[59] R. Penrose, Applications of negative dimensional tensors, Comb. Math. Appl. **1**, 221 (1971).

[60] N. Schuch, D. Pérez-García, and I. Cirac, Classifying quantum phases using matrix product states and projected entangled pair states, Phys. Rev. B **84**, 165139 (2011).

[61] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete, Matrix product states and projected entangled pair states: Concepts, symmetries, theorems, Rev. Mod. Phys. **93**, 045003 (2021).

[62] G. M. Crosswhite, A. C. Doherty, and G. Vidal, Applying matrix product operators to model systems with long-range interactions, Phys. Rev. B **78**, 035116 (2008).

[63] G. M. Crosswhite and D. Bacon, Finite automata for caching in matrix product algorithms, Phys. Rev. A **78**, 012356 (2008).

[64] S. Paeckel, T. Köhler, and S. R. Manmana, Automated construction of $U(1)$-invariant matrix-product operators from graph representations, SciPost Phys. **3**, 035 (2017).

[65] B. Bauer, P. Corboz, R. Orús, and M. Troyer, Implementing global Abelian symmetries in projected entangled-pair state algorithms, Phys. Rev. B **83**, 125106 (2011).

[66] S. Singh, R. N. C. Pfeifer, and G. Vidal, Tensor network states and algorithms in the presence of a global U(1) symmetry, Phys. Rev. B **83**, 115125 (2011).

[67] S. Singh and G. Vidal, Tensor network states and algorithms in the presence of a global SU(2) symmetry, Phys. Rev. B **86**, 195114 (2012).

[68] A. Weichselbaum, Non-Abelian symmetries in tensor networks: A quantum symmetry space approach, Ann. Phys. **327**, 2972 (2012).

[69] A. Weichselbaum, X-symbols for non-Abelian symmetries in tensor networks, Phys. Rev. Res. **2**, 023385 (2020).

[70] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, Matrix Product Density Operators: Simulation of Finite-Temperature and Dissipative Systems, Phys. Rev. Lett. **93**, 207204 (2004).

[71] I. Pižorn, L. Wang, and F. Verstraete, Time evolution of projected entangled pair states in the single-layer picture, Phys. Rev. A **83**, 052321 (2011).

[72] H. C. Jiang, Z. Y. Weng, and T. Xiang, Accurate Determination of Tensor Network State of Quantum Lattice Models in Two Dimensions, Phys. Rev. Lett. **101**, 090603 (2008).

[73] E. R. Davidson, The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices, J. Comput. Phys. **17**, 87 (1975).

[74] H. N. Phien, J. A. Bengua, H. D. Tuan, P. Corboz, and R. Orús, Infinite projected entangled pair states algorithm improved: Fast full update and gauge fixing, Phys. Rev. B **92**, 035142 (2015).

[75] G. Evenbly, Gauge fixing, canonical forms, and optimal truncations in tensor networks with closed loops, Phys. Rev. B **98**, 085155 (2018).

[76] A. Buckley and A. Lenir, Algorithm 630: BBVSCG–a variable-storage algorithm for function minimization, ACM Trans. Math. Softw. **11**, 103 (1985).

[77] D. C. Liu and J. Nocedal, On the limited memory BFGS method for large scale optimization, Math. Program. **45**, 503 (1989).

[78] This code was developed for momentum-space and hybrid momentum-real space calculations [32–34], but is also quite efficient for pure real-space calculations.

[79] P. Corboz, Improved energy extrapolation with infinite projected entangled-pair states applied to the two-dimensional Hubbard model, Phys. Rev. B **93**, 045116 (2016).