

Algorithms for finite projected entangled pair states

Michael Lubasch, J. Ignacio Cirac, and Mari-Carmen Bañuls

Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Straße 1, 85748 Garching, Germany

(Received 15 May 2014; revised manuscript received 8 August 2014; published 25 August 2014)

Projected entangled pair states (PEPS) are a promising ansatz for the study of strongly correlated quantum many-body systems in two dimensions. However, due to their high computational cost, developing and improving PEPS algorithms is necessary to make the ansatz widely usable in practice. Here we analyze several algorithmic aspects of the method. On the one hand, we quantify the connection between the correlation length of the PEPS and the accuracy of its approximate contraction and discuss how purifications can be used in the latter. On the other hand, we present algorithmic improvements for the update of the tensor that introduce drastic gains in the numerical conditioning and the efficiency of the algorithms. Finally, the state-of-the-art general PEPS code is benchmarked with the Heisenberg and quantum Ising models on lattices of up to 21×21 sites.

DOI: [10.1103/PhysRevB.90.064425](https://doi.org/10.1103/PhysRevB.90.064425)

PACS number(s): 02.70.-c, 75.10.Jm, 64.70.Tg, 03.67.-a

I. INTRODUCTION

As an ansatz for the wave function of a quantum many-body system, projected entangled pair states (PEPS) [1] represent the natural generalization of matrix product states (MPS) to higher dimensions. MPS are by now well established as a numerical tool; they constitute the variational class on which the density matrix renormalization group (DMRG) is based upon [2,3] and, nowadays, DMRG is considered numerically exact for systems comprising hundreds of quantum particles in one dimension.

PEPS have the potential to reproduce this success in higher dimensions. This is particularly interesting for problems that cannot be tackled otherwise, e.g., two-dimensional fermionic or frustrated systems where quantum Monte Carlo methods are hampered by the notorious sign problem. First PEPS results for such problems in condensed matter as the t - J model or the Kagome antiferromagnet compare well to the best currently known results achieved by other means [4–6]. However, the significantly higher computational cost of PEPS algorithms compared to MPS restricts the feasible simulations to PEPS tensors with much smaller dimensions, and the results are still far from the level of convergence attained in the one-dimensional case.

In the last years, significant conceptual and algorithmic progress has been made, e.g., Refs. [7–15]. Many of the numerical studies have focused on systems in the thermodynamic limit, for which the iPEPS [9] ansatz can be used. In such case, the translational invariance of the system is exploited to reduce the number of variational parameters to the few tensors in a small unit cell. However, the nontranslationally invariant finite PEPS ansatz is also of great importance. On the one hand, by avoiding a predefined unit cell it allows a more unbiased approach to the thermodynamic limit, when combined with finite size scaling (although, also, a systematically increased unit cell in iPEPS can be expected to produce more and more unbiased results [5,6]). On the other, it is the proper ansatz for problems that are intrinsically nontranslationally invariant, such as the simulation of current optical lattice experiments that are being carried out in inhomogeneous traps. In exchange, the price to pay is a more involved implementation and longer running times that scale with the system size.

The original PEPS algorithms [1,16] can cope with the nontranslationally invariant situations [17–20], but a

straightforward implementation attains only small tensor dimensions, and is not enough to explore the power of the ansatz. In order to reach larger dimensions (i.e., comparable to those used in present iPEPS calculations) and to approach the optimal ground-state approximations for them, it is necessary to take into account and to optimize the cost and stability of every stage of the algorithms, which is only feasible through a thorough understanding of the various possibilities. Only then it will be possible to adopt the optimal strategies for the particular problem at hand.

In this paper, we aim at a global understanding of the algorithmic aspects of finite PEPS, both at the physical and technical level. We address the two fundamental ingredients of PEPS algorithms, namely the environment approximation, i.e., the approximate contraction of the tensor network (TN), and the tensor update. In a previous work [21], we focused on the environment approximation. We studied the physical significance and limitations of various contraction strategies, and introduced the cluster scheme, which unifies previous methods and gives rise to a new contraction algorithm with a trade-off between precision and computational cost. Here, we extend the analysis of the environment approximation to provide new insight into the convergence of the cluster strategy by relating it to the correlation length of the system. Additionally, we show how the environment approximation can be kept exactly positive with the help of purifications. Regarding the tensor update, we investigate the effect of restricting the variational parameters to a reduced tensor [12], a technique used often in the case of iPEPS and characterized by a lower computational cost. Via the reduced tensor we derive new numerical methods, namely suitable gauge choices, which significantly enhance the stability of the update algorithm. These gauge choices admit also a generalization to cases where the full tensor needs to be updated.

Furthermore, we benchmark the state-of-the-art finite PEPS algorithms using the Heisenberg Hamiltonian and the quantum Ising model with transverse field. By presenting converged finite PEPS results for lattice sizes typically considered in the context of finite size scaling, we not only assess the validity of the ansatz, but enable a systematic comparison to other methods and implementations.

The rest of the article is structured as follows. In Sec. II, we briefly present the basic notation and concepts common to PEPS algorithms. The algorithmic details regarding the

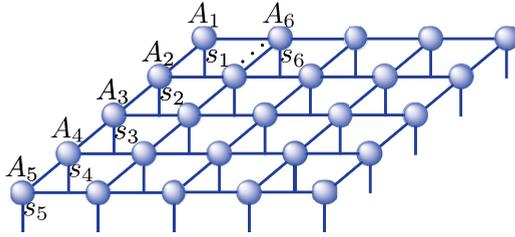


FIG. 1. (Color online) A PEPS on a 5×5 square lattice. Physical indices are depicted pointing down, and virtual ones connect tensors in the plane.

convergence of the cluster scheme, the use of positive environments and the strategies to improve the tensor update are discussed in Sec. III. Section IV collects the numerical results corresponding to our best PEPS ground-state approximations for the benchmark models. Finally, in Sec. V, we summarize our conclusions.

II. NOTATION AND PRELIMINARY CONCEPTS

For completeness, we introduce here the main concepts that will be used throughout this paper. Reviews on PEPS and more general TN methods can be found in the literature [16,22].

Given a quantum system of N particles, with dimensions d_l and local Hilbert space bases $\{|s_l\rangle\}_{s_l=1}^{d_l}$ (for $l = 1, \dots, N$), a PEPS [1] is a state of the form

$$|\psi^{\text{PEPS}}\rangle := \sum_{s_1, s_2, \dots, s_N} \mathcal{F}(A_1^{s_1} A_2^{s_2} \dots A_N^{s_N}) |s_1 s_2 \dots s_N\rangle,$$

where \mathcal{F} denotes the contraction of a TN formed by the tensors $A_l^{s_l}$. Figure 1 shows the two-dimensional square lattice geometry with open boundary conditions and size $N = L \times L$ considered throughout this work. In this geometry, to each lattice site l corresponds a tensor $A_l^{s_l}$ with one physical index s_l for its physical degree of freedom and up to four virtual indices connecting neighboring tensors. The dimension of the virtual indices, called *bond dimension* D , restricts the maximum possible block entropy of the state according to an area law.

PEPS algorithms for finding ground states can be classified in two types, namely, variational minimization of energy and imaginary time evolution. With only minor changes, the second one allows also the simulation of real time evolution. Both kinds of algorithm can be formulated in terms of the minimization of a certain cost function by varying the tensor parameters.¹ This minimization is realized in practice by means of an alternating least squares (ALS) scheme, in which one sweeps over the tensors and updates them one after another, each time choosing the components that minimize the cost function under the constraint that all the other tensors are fixed.

Throughout this paper we focus (almost exclusively) on the imaginary time evolution. In this case, it is customary to use a

¹We remark that the initial time-evolution method [23–25] works slightly different, and the basis of this work is rather given by Ref. [26].

Suzuki-Trotter approximation of the evolution operator where the Hamiltonian is split into parts containing only mutually commuting terms. The cost function to be minimized is then the distance $d(|\psi\rangle) = \|\psi - \hat{O}|\phi\rangle\|^2$, where $|\phi\rangle$ is the initial PEPS, \hat{O} is an operator representing one (or more) Trotter gates, i.e., the exponential of one (or several) such Hamiltonian terms [16], and $|\psi\rangle$ is the resulting PEPS. During the ALS sweeping, the tensor for site l is the one that minimizes

$$d(A_l) =: \vec{A}_l^\dagger N_l \vec{A}_l - \vec{A}_l^\dagger \vec{b}_l - \vec{b}_l^\dagger \vec{A}_l + \text{const.} \quad (1)$$

It is given by the solution of the linear system of equations $N_l \vec{A}_l = \vec{b}_l$, i.e., $\vec{A}_l = N_l^{-1} \vec{b}_l$, where the norm matrix N_l results from the norm TN $\langle\psi|\psi\rangle$ by leaving out the tensor A_l^* in the bra and A_l in the ket, and \vec{b}_l results from the TN $\langle\psi|\hat{O}|\phi\rangle$ by leaving out A_l^* in the bra. This procedure can be iterated for the necessary number of steps to reach the desired total (real or imaginary) time.²

Two main parts, namely, the *environment approximation* and the *tensor update*, constitute the building blocks of this algorithm and will be often referred to in the rest of this article. The first notion corresponds to the exact or approximate evaluation of the effective matrix (N_l) and vector (\vec{b}_l) that determine the local equation to be solved for the tensor at a given site. The second term denotes the solution of the vector equation and the corresponding change of the PEPS with the updated tensor.

Some strategies developed in the context of iPEPS can also be applied to the finite case, and we will do that in the following. The most widely used iPEPS method, due to its efficiency and stability, is the simple update (SU) [8], in which the environment is assumed to be separable and then the tensors are updated via simple SVD. As we showed in Ref. [21], the SU works equally with finite PEPS but produces results with limited accuracy. The full update (FU) [9,12] is based on a more accurate approximation of the environment, in closer analogy to the original finite PEPS algorithm [1,16], but differing from it in the fact that Trotter gates are not applied simultaneously, so that the environment for the update of one gate does only require the norm contraction around that gate. We will in the following use the term FU in the context of finite PEPS to denote the sequential application of Trotter gates together with the full contraction of the norm TN, as in Ref. [21].

III. ALGORITHMIC ASPECTS

In this section we analyze several distinct aspects of finite PEPS algorithms, regarding both the environment approximation and the tensor update.

²In the direct minimization of the energy, the cost function to be minimized is $E(|\psi\rangle) = \langle\psi|\hat{H}|\psi\rangle/\langle\psi|\psi\rangle$ over the PEPS $|\psi\rangle$. For an update during the ALS sweeping, the tensor at position l is set to A_l minimizing $E(\vec{A}_l) = \vec{A}_l^\dagger H_l \vec{A}_l / \vec{A}_l^\dagger N_l \vec{A}_l$, which can be found as the lowest eigenstate of the generalized eigenproblem $H_l \vec{A}_l = \lambda N_l \vec{A}_l$. Here, the matrix H_l is defined from the TN $\langle\psi|\hat{H}|\psi\rangle$, in which one leaves out the tensor A_l in the ket and A_l^* in the bra, and N_l is the norm matrix.

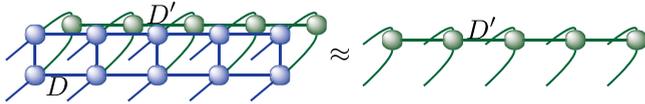


FIG. 2. (Color online) Original contraction [1,16] of the norm TN for a PEPS with bond dimension D . The product of a bulk row with a boundary MPO of bond dimension D' is approximated by a new boundary MPO of the same dimension, at a cost $\mathcal{O}(dD^6D^2) + \mathcal{O}(D^4D^3)$.

In particular, for the environment approximation we show how the success of the cluster scheme introduced in Ref. [21] is deeply connected to the correlation length of the state. We also discuss the feasibility and the cost of explicitly keeping a positive environment by making use of purification MPO.

For the tensor update, we propose gauge choices for each possible update scheme, and show how they improve the numerical stability of the algorithms. We additionally discuss how the reduced tensor, originally introduced in the context of iPEPS [12], can similarly be used in the finite case to speed up the computations. The normalization of the tensors is another factor that can improve the stability of the method.

While part of this section is significantly technical, the considerations exposed here are relevant for the implementation of any (finite) PEPS algorithm. Furthermore they have also clear physical implications, especially in the case of the environment contraction.

A. Environment approximation

In the imaginary time algorithm, the update of one tensor at lattice site l involves the contraction around that tensor of the norm TN $\langle \psi | \psi \rangle$ and of the TN $\langle \psi | e^{-\tau \hat{H}_x} | \phi \rangle$ for a certain subset of (mutually commuting) Hamiltonian parts \hat{H}_x . The first contraction leads to the norm matrix N_l and the second to the vector \vec{b}_l from which the new tensor for that lattice site follows as $\vec{A}_l = N_l^{-1} \vec{b}_l$. The original algorithm [1,16] includes the complete $e^{-\tau \hat{H}_x}$, i.e., all (mutually commuting) Trotter gates, in the TN $\langle \psi | e^{-\tau \hat{H}_x} | \phi \rangle$ and thus requires two independent environment approximations, one for $\langle \psi | \psi \rangle$ and one for $\langle \psi | e^{-\tau \hat{H}_x} | \phi \rangle$. However, in the following we adopt the strategy from Ref. [21]. If Trotter gates are applied one by one, and only the tensors on which a given gate acts are modified, then it suffices to consider the environment approximation of the norm TN alone, and, starting from this environment, the vector \vec{b}_l is constructed from the exact contraction of a single Trotter gate.³

As in the original algorithm [1,16], we can approximate the environment of a PEPS row (column) with the help of *boundary MPO*. By identifying two opposite sides of the PEPS TN with boundary MPO, the action of intermediate rows (columns) on those is successively approximated by new boundary MPO, as shown in Fig. 2 for the norm, until the tensors of interest are reached. The approximation accuracy

³In an efficient implementation of this algorithm, we only store and update the boundary MPO for the rows and columns of the norm TN, and the contraction for the vector \vec{b}_l is performed on the fly.

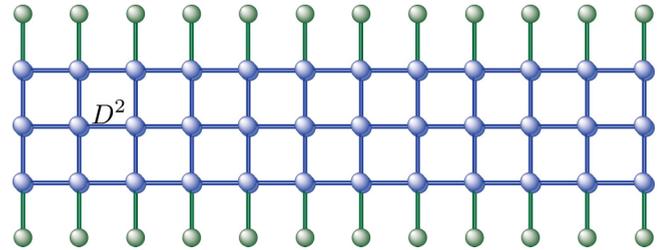


FIG. 3. (Color online) Cluster environment of a bulk row for cluster size $\delta = 1$. The contraction outside the cluster is achieved by means of a separable positive boundary MPO in the way explained in Ref. [21] with computational cost $\mathcal{O}(dD^5)$. Then the cluster is contracted with the help of a general boundary MPO of bond dimension D' , which, in this case, can be found with $\mathcal{O}(dD^5D^2)$ operations.

of this method, used also in the FU and the cluster scheme, is ultimately determined by the *boundary bond dimension* D' of the boundary MPO, and its efficiency is dictated by the leading computational cost $\mathcal{O}(dD^6D^2) + \mathcal{O}(D^4D^3)$. In typical calculations, the boundary bond dimension for a certain approximation precision scales as $D' \propto D^2$ independent of the system size, such that the original contraction [1,16] has the overall cost $\mathcal{O}(D^{10})$.

1. Accuracy of the cluster contraction

The cluster update (CU) introduced in Ref. [21] allows a trade-off between precision and efficiency in the environment approximation. The SU and the FU are special cases of this procedure, which naturally interpolates between them in both accuracy and computational cost. Because clusters are not only useful for the tensor update but equally for the computation of expectation values, they realize a unifying framework for PEPS contractions.

A cluster is defined as a set of tensors comprising the considered ones and their neighborhood up to a distance called *cluster size* δ . The idea is to approximate the environment outside the cluster very roughly and inside with more precision. In the context of finite PEPS algorithms, in which the TN is contracted row by row, it is reasonable to define a cluster as the considered row and its neighboring rows up to the distance δ . Figure 3 shows an example cluster of size $\delta = 1$ around a central row in the bulk of a PEPS. We employ a separable positive boundary MPO for the contraction outside the cluster, which has a cost $\mathcal{O}(dD^5)$, while the cluster itself is contracted using a general boundary MPO of bond dimension $D' > 1$, which in this case requires $\mathcal{O}(dD^5D^2)$ operations [21]. The precision and efficiency of the approximation are determined by the cluster size and D' . The separable positive boundary MPO produces for cluster size $\delta = 0$ an environment approximation equivalent to the SU one [21].

We previously observed [21] that the contraction error decreases exponentially with the cluster size for PEPS ground-state approximations of the Heisenberg model. This property, which justifies the usability of clusters, is ultimately related to a finite correlation length of the system, as we appreciate here with the help of the quantum Ising model. This model becomes critical in the thermodynamic limit at transverse field

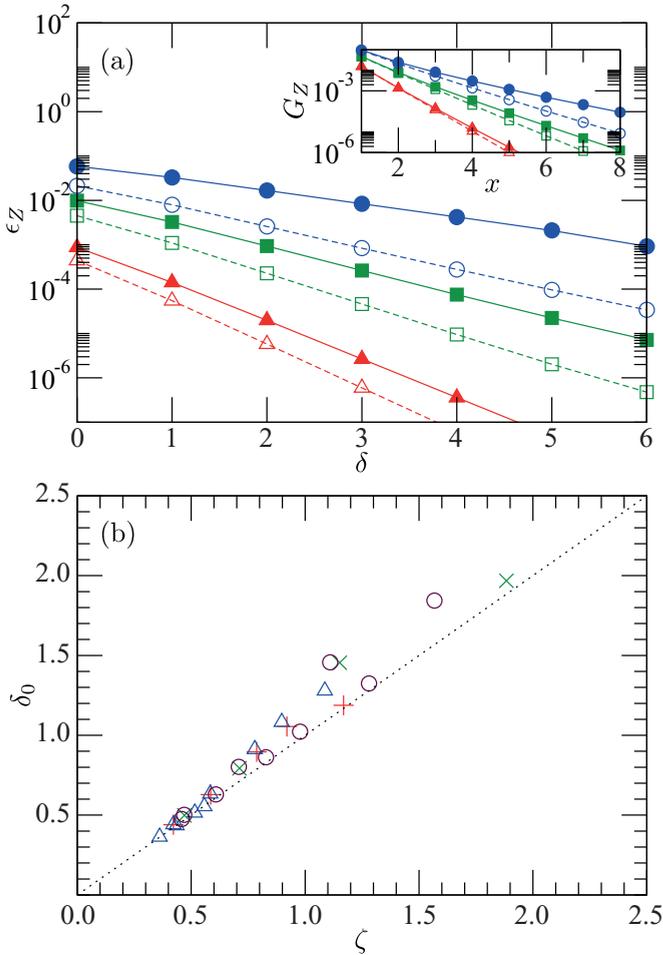


FIG. 4. (Color online) Relation between cluster error and correlation function in the Ising model on a 21×21 lattice. (a) Cluster error (main plot) and correlation function (inset) for observable σ^Z , for $D = 2$ (open symbols), 3 (filled symbols), and $B = 2.0$ (triangles), 2.5 (squares), 2.8 (circles). (b) Characteristic cluster size, δ_0 , vs correlation length, ζ , for several values of $B \in [2, 4]$, for observable σ^Z with $D = 2$ (pluses), 3 (crosses), and for σ^X with $D = 2$ (triangles), 3 (circles).

$B \approx 3.044$, and, thus, by varying B we can create states with different correlation lengths.

We have analyzed the cluster contraction error of a local observable acting on the center of the lattice, $\epsilon_\alpha(\delta) := |\langle \sigma^\alpha \rangle_\delta - \langle \sigma^\alpha \rangle| / |\langle \sigma^\alpha \rangle|$, for $\alpha = X, Z$, where $\langle \sigma^\alpha \rangle_\delta$ is the approximated contraction using cluster size δ , and $\langle \sigma^\alpha \rangle$ the result of contracting the full TN. The behavior of this quantity can be compared to the correlation function, $G_\alpha(x) := \langle \sigma_l^\alpha \sigma_{l+x}^\alpha \rangle - \langle \sigma_l^\alpha \rangle \langle \sigma_{l+x}^\alpha \rangle$, for two sites separated by a distance x along the central column of the lattice. All contractions were performed with large enough $D' = 100$ such that the contraction error was independent of D' .⁴ We observe in Fig. 4(a) that the decrease of the contraction error is always steeper for a faster decaying correlation function. In order to make this statement more

⁴The dependence of the cluster contraction on D' was already investigated previously [21].

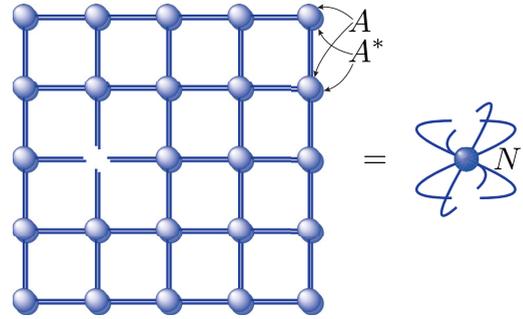


FIG. 5. (Color online) Norm environment for a single site of a 5×5 PEPS. Because each PEPS tensor A from Fig. 1 is contracted with its complex conjugate A^* over the physical index, an exact contraction of this norm TN would give a positive Hermitian norm matrix N . For large PEPS, an exact contraction of N is not feasible, and the original contraction approximation [1,16] based on general boundary MPO, shown in Fig. 2, does not keep the positivity.

precise, we can fit the decay of the error to an exponential function of the cluster size, $\epsilon_\alpha(\delta) \propto \exp(-\delta/\delta_0)$, and obtain a characteristic cluster size δ_0 . Correspondingly, we can extract a correlation length ζ from a similar fit of the correlation function $G_\alpha(x) \propto \exp(-x/\zeta)$. After having calculated δ_0 and ζ for several PEPS,⁵ we plot δ_0 as a function of ζ in Fig. 4(b) and conclude that $\delta_0 \approx \zeta$.⁶ This demonstrates an extremely clear quantitative connection between the cluster contraction error for a given cluster size and the correlations in the state.

2. Positive environment

The exact norm environment, resulting from an exact contraction of $\langle \psi | \psi \rangle$ around one (or several) site(s), is positive by construction, as can be seen in Fig. 5. Although this positive characteristic is considered a desirable property for the environment approximation, in general it is not respected by the approximated contractions. Nevertheless, it is possible to use schemes that maintain it. In particular, the single-layer (SL) algorithm was introduced in Ref. [14] as a way to improve the environment approximation of the SU while preserving its efficiency and numerical stability. The SL method performs the norm contraction by means of transformations in the ket alone [14]. Then the boundary is described by a *purification* MPO [26], defined via a MPS of virtual bond dimension D'' and physical dimension $D \times d'$ in such a way that the MPO results from tracing over the *purification bonds* of dimension d' . Approximating the environment in the SL way and then updating the tensors as explained in Ref. [14] ensures a stable algorithm, but as seen in Ref. [21] the error in the environment approximation can be several orders of magnitude above that

⁵We determine δ_0 via the two values of $\epsilon_\alpha(\delta)$ at $\delta = 2$ and 4, and ζ via the two values of $G_\alpha(x)$ at $x = 4$ and 8.

⁶The slight deviation of Fig. 4(b) from the exact diagonal for larger ζ is due to the fact that they correspond to $B \approx 3.0$ where both $\epsilon_\alpha(\delta)$ and $G_\alpha(x)$ decrease rather polynomially and thus δ_0 as well as ζ depend more strongly on the two values from which they were determined.

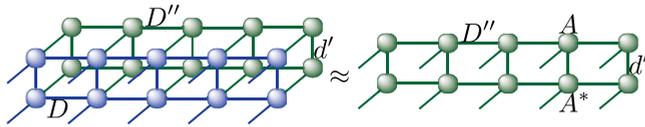


FIG. 6. (Color online) Positive contraction of the norm TN for a PEPS with bond dimension D . The product of a bulk row with a boundary purification of virtual bond dimension D'' and purification bond dimension d' is approximated by a new boundary purification of the same dimensions. Each update of a tensor A constitutes a nonlinear problem, and solving the linearized equations costs $\mathcal{O}(dD^6D''^4) + \mathcal{O}(D^4D''^6) + \mathcal{O}(d^3D^3D''^6)$ where $d' \leq DD''^2$.

of the original contraction [1,16], and it can depend strongly on the system size.

Several factors can cause these accuracy limitations. Even if there exists a good positive MPO approximation for the boundary with moderate bond dimension D' (as observed for gapped systems [27]), it does not necessarily follow that D'' is small [28,29] and hence it is not clear a priori that fixing the maximum D'' produces an accurate approximation for the environment. Moreover, as argued in Ref. [21], by operating on a single layer, the scheme does not find the most general purification with given bond D'' . Here we want to address the question whether the accuracy limitations of the SL algorithm are due to the description of the boundary as purification or whether they are due to the specific operations proposed in Ref. [14] to determine that boundary purification.

One way to allow for a more general purification is to formulate an algorithm in the double-layer picture in the following way. Given the MPS bond dimension, D'' , and the purification bond, d' , we write a purification MPO [26] by using in the lower layer the complex conjugated tensors from the upper layer. The problem of approximating the boundary after the contraction of one further row of the PEPS norm TN is then formulated for this structure instead of the general MPO, as sketched in Fig. 6. The local equations result from replacing the single tensor of the general MPO by the structure consisting of A_l and A_l^* . Following the standard ALS procedure, we sweep over the sites l , and for each site solve the corresponding optimization problem for A_l . However, in this case the cost function to be minimized is no longer quadratic, but quartic in the variables of a tensor at site l , and its minimum corresponds to the solution of nonlinear equations, in contrast to the linear equations encountered in the original contraction [1,16] of Fig. 2. The nonlinear equations for A_l have to be solved iteratively. We describe and benchmark several options in Appendix A.

We compare this scheme to the original [1,16] and the SL algorithm based on the norm contraction of the same PEPS used in the previous analysis of Ref. [21]. The results are shown in Fig. 7. For a fixed purification bond d' , we observe that the relative error of the norm decreases fast as a function of D'' . The comparison of the 11×11 to the 21×21 lattice [Fig. 7(a)] shows that, similar to the original algorithm [1,16], the error does not have a strong dependence on the system size. On the other hand, with growing d' , the curves tend to converge to the error of the original contraction [1,16]. This effect can be observed already with small purification bonds for

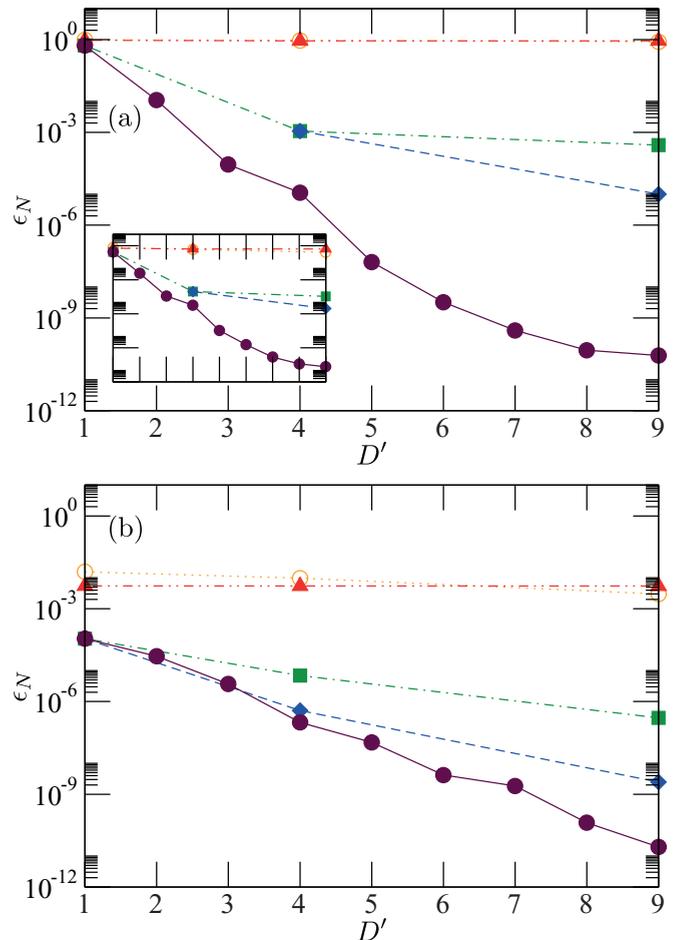


FIG. 7. (Color online) Relative error of the norm contraction using boundary purification MPO, for SU Ising ground-state approximations. (a) $B = 3.0$ and $D = 2$ on lattices of size 21×21 (main plot) and 11×11 (inset). (b) $B = 1.0$ and $D = 4$ on a 11×11 lattice. For reference, we show the error of the SL method with maximum $d' = DD''^2$ (open circles) and of the original algorithm [1,16] (filled circles). Our purification contraction was performed with $d' = 1$ (triangles), 2 (squares), and 3 (diamonds), and $D'' = \sqrt{D'}$.

$B = 1.0$ [Fig. 7(b)]. These results suggest that the error in the SL method is mainly due to the restricted class of purifications it can attain, and not to the description of the boundary as a purification with small bond D'' .⁷

From the discussion above, we conclude that it is possible to efficiently find a (close to) optimal general purification by means of the solution of nonlinear equations. In the context of PEPS contractions, this technique improves the SL scheme significantly, but given its higher computational cost compared to the original contraction [1,16], resulting from the iterative

⁷Although our conclusions are based on PEPS from the SU, because we had thoroughly analyzed exactly the same PEPS with the SL algorithm in Ref. [21], we clearly expect similar improvements for PEPS from the FU (in particular the ground state of the quantum Ising model at $B = 1.0$, away from criticality, should be well approximated by the SU and thus Fig. 7(b) should only change slightly for the corresponding FU PEPS).

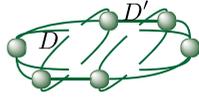


FIG. 8. (Color online) Six environment tensors of a nearest-neighbor tensor pair. They form a periodic boundary MPO with virtual bond dimension D' and physical dimension D . It is constructed with $\mathcal{O}(dD^6D'^2) + \mathcal{O}(D^4D'^3)$ operations, resulting from the optimal search for a boundary MPO and the contraction of the environment up to the tensor pair.

routines (see Appendix A), it is not a practical option. Hence, in the following, all our cluster and full contractions will be based upon the original contraction algorithm [1,16] and thus make use of general boundary MPO as shown in Fig. 2. Nevertheless, the procedures analyzed here may be useful for other problems where the question of numerically optimizing a purification MPO appears, such as for the description of one-dimensional thermal states or open systems.

B. Tensor update

Once the environment is computed, the actual update of the tensors takes place by solving the appropriate local equations. It is also possible to use simplifications of this step which render a more efficient and stable algorithm.

For the update of a pair of neighboring tensors, the environment can be approximated in general by a MPO with periodic boundary conditions, as illustrated in Fig. 8. A first simplification of the tensor update procedure comes from sequentially processing the Trotter gates, as described above, and changing only the tensors on which each gate acts. Then all the update operations on the pair are performed with a fixed environment.

The computational cost of the tensor update can be greatly decreased by restricting the update to the *reduced tensor* [12]. This reduced tensor update minimizes the cost function (1) for the full tensors exactly only in the case of a separable environment [21] as, e.g., in the SU, but it is worth studying its performance in a more general situation. In any case, it allows to work with larger bond dimensions, which might compensate for the smaller number of variational parameters.

Another major difference between MPS and PEPS concerns the conditioning of the effective norm matrix N_l . For MPS with open boundary conditions, a *gauge transformation*⁸ can be chosen such that $N_l = 1$, which guarantees the stability of the tensor update. Although this is impossible for PEPS, we will show how a proper gauge choice and tensor normalization drastically improve the stability of the algorithm.

⁸In a PEPS, for any pair of neighboring tensors that are connected via a virtual index, an arbitrary matrix M can be contracted with one tensor and the matrix M^{-1} with the other in such a way that the state does not change. This establishes the analog of a (local) gauge freedom. Throughout this article, the matrix M is called *gauge matrix* or *gauge transformation* in accordance with Ref. [16], and the term *gauge fixing* or *gauge choice* refers to the process of choosing a specific matrix M .

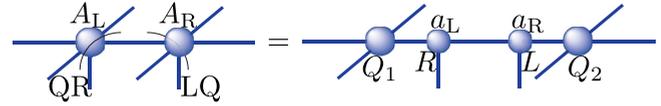


FIG. 9. (Color online) A QR decomposition of the left full tensor A_L generates the left reduced tensor a_L as the R . Similarly, a LQ decomposition of the right full tensor A_R gives the right reduced tensor a_R as the L . The initial dD^4 variational parameters of the full tensor are decreased to the d^2D^2 variational parameters of the reduced tensor.

1. Reduced tensor

Before performing the update under a nearest-neighbor Trotter gate, the tensor for a lattice site can be decomposed into the contraction of two tensors, in such a way that one of them carries the physical index and the virtual bond corresponding to the link on which the two-site gate acts, i.e., all the indices directly affected by the gate. This tensor is called the *reduced tensor* [12], and can be obtained from the full tensor by means of a QR decomposition, as sketched in Fig. 9.

In the reduced tensor update, only the components of such reduced tensor are modified during the update procedure, while the remaining part of the full tensor is left unchanged. These remaining parts of both tensors in the pair are contracted with the periodic MPO of Fig. 8 to get the environment for the reduced tensor pair, N_{red} , shown in Fig. 10. Due to the approximate contractions, this reduced environment is in general not positive, neither is it Hermitian, but its positive approximant can be constructed in two steps [30]. First, we compute the optimal Hermitian approximant $\tilde{N}_{\text{red}} := (N_{\text{red}} + N_{\text{red}}^\dagger)/2$. Second, from its eigendecomposition $\tilde{N}_{\text{red}} = U\Sigma U^\dagger$ we obtain the positive approximant as $U\Sigma_+U^\dagger$, where Σ_+ results from Σ by setting all negative eigenvalues to zero. Finally, the environment is written as $\tilde{X}\tilde{X}^\dagger$ in terms of its square root $\tilde{X} := U\sqrt{\Sigma_+}$.

The computational cost of the contractions for the periodic boundary MPO (Fig. 8), needed in both the reduced and the full tensor update, reads $\mathcal{O}(dD^6D'^2) + \mathcal{O}(D^4D'^3)$. The construction of N_{red} (Figs. 9 and 10) is only slightly more expensive with $\mathcal{O}(d^4D^4D'^2) + \mathcal{O}(d^2D^6D'^2) + \mathcal{O}(d^2D^4D'^3)$ operations. Its eigendecomposition requires $\mathcal{O}(d^6D^6)$. In the complete update of the reduced tensors via the sweeping of the ALS scheme, all further operations have lower computational cost. Notice that in the case of the full tensors, the contraction of the norm environment for a single tensor needs

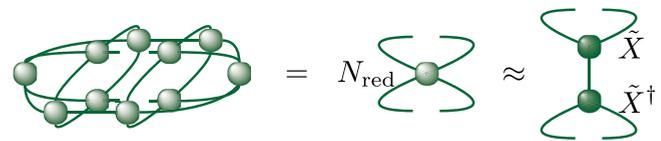


FIG. 10. (Color online) Environment tensor N_{red} of a reduced nearest-neighbor tensor pair and its closest positive semidefinite approximant $\tilde{X}\tilde{X}^\dagger$ constructed as explained in the text. The contractions are characterized by the leading computational cost $\mathcal{O}(d^4D^4D'^2) + \mathcal{O}(d^2D^6D'^2) + \mathcal{O}(d^2D^4D'^3)$ and the computation of the positive approximant requires additionally $\mathcal{O}(d^6D^6)$ operations.

$\mathcal{O}(D^8 D^2) + \mathcal{O}(D^4 D^3)$ operations while the eigendecomposition of the norm matrix has the cost $\mathcal{O}(d^3 D^{12})$.

In order to study whether the reduced tensor limits the accuracy of the method, we considered imaginary time evolution of the Heisenberg model on 4×4 and 10×10 lattices, and compared the final energies from the reduced tensor update to the ones from the full tensor update. We found that, while for the small bond dimensions $D = 2$ and 3 the full tensor update produced better energies, for $D = 4$ the energies of both approaches were already very similar. This can be appreciated by comparison of the results in Appendix C (obtained with the reduced tensor) to the full tensor results published in Refs. [1,18].

Because the reduced tensor update is less costly, we could reach larger bond dimensions than with the full tensor update, and, in the end, obtained the lowest energies with the reduced tensors. Therefore the reduced tensor update was used for the results presented in this paper.

2. Gauge fixing

In the case of MPS, it is possible to keep up a canonical form of the tensors during their updates with the help of a local gauge fixing, and that ensures the stability of the algorithm and optimizes its performance [16]. In the case of PEPS, there exists neither such a canonical form nor any means to locally gauge away the norm matrix. Nevertheless, using the gauge freedom, it is possible to improve the conditioning of the norm matrix and positively affect the precision and stability of the method, as we describe in the following.⁹

We propose a gauge fixing that is inspired by the one-dimensional case with open boundary conditions. In that case, the norm tensor can be reduced to the identity by (partially) imposing the canonical form of the MPS, achieved by QR (or LQ) decomposition of each tensor after its update [16]. Alternatively, for an arbitrary MPS, it is always possible to reduce the norm matrix to the identity by taking the square roots of the unconnected left and right environment halves and absorbing part of their QR (LQ) decompositions in the tensor to be updated.

In the case of PEPS, it is not possible to ensure an identity norm matrix by means of QR or LQ decompositions after the tensor update. Hence we adapt the second possibility and obtain the gauge transformations from the environment before the tensor update, namely from the norm tensor itself, such that the norm matrix is better conditioned. Because this gauge fixing can be combined with any of the environment approximations described previously, we propose a precise scheme for each case.

When the environment of the tensor pair is separable, i.e., $D' = 1$ in Fig. 8, it decomposes into six positive semidefinite

⁹A recent alternative approach for an improved imaginary time evolution is presented in Ref. [31], in which the authors propose a “quasicanonical form” that arises as fix point of the SU performed with nearest-neighbor identity gates, and in which they demonstrate how that form allows an efficient and stable imaginary time evolution with projected entangled pair operators (PEPO) in the context of iPEPS.

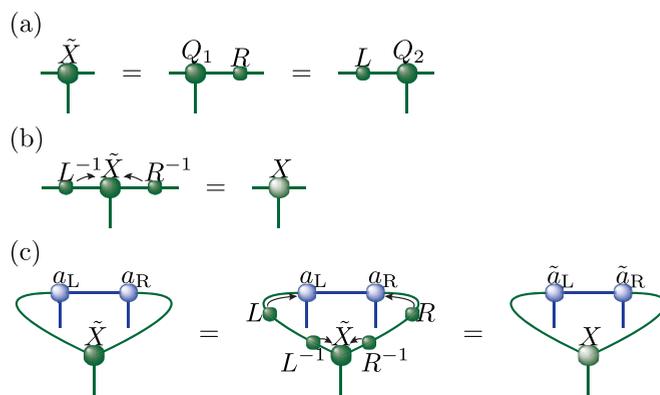


FIG. 11. (Color online) Gauge fixing on the environment tensor of the reduced tensor pair, when the environment is nonseparable. (a) We perform a QR and LQ decomposition on \tilde{X} from Fig. 10 independently of each other (notice that we have shortened here the horizontal open indices of \tilde{X} compared to Fig. 10). (b) Contraction of \tilde{X} with L^{-1} and R^{-1} gives the final square root of the environment tensor X . (c) In order to leave the state unchanged, the left and right reduced tensors a_L and a_R from Fig. 9 have to be contracted with the gauge transformations L and R as shown here, which gives the starting tensors \tilde{a}_L and \tilde{a}_R for the update explained in Fig. 12.

matrices, which can be determined by the algorithm in Ref. [21]. We compute the square roots of these matrices and absorb them in the tensor pair. After contraction of the tensor pair with the Trotter gate, a SVD is performed to find the new tensors, and finally these are multiplied by the inverses of the previous square roots. This procedure coincides with the SU [8] in which the λ matrices surrounding the tensor pair are substituted here by the square roots of the environment matrices corresponding to each link. Since the positive separable environment of the tensor pair is obtained with $\mathcal{O}(dD^5)$ operations, the leading cost of the complete update is $\mathcal{O}(d^6 D^3) + \mathcal{O}(d^2 D^5)$, under the assumption $d \leq D^2$.

When the environment of the tensor pair is nonseparable, and we restrict the update to the reduced tensor, we propose the gauge fixing from Fig. 11. By taking R and L from independent QR and LQ decompositions of the same \tilde{X} from Fig. 10, we treat both virtual bonds of the environment equally, such that both reduced tensors will experience similar condition numbers in the linear equations of the following sweeping.¹⁰ After we have obtained the desired better-conditioned square root of the environment tensor, X , in order to leave the state unchanged, the left (right) reduced tensor has to be contracted with L (R) over its left (right) virtual index.

After our gauge fixing has been applied, the actual update takes place in three steps. First, the tensors are initialized using a SVD as shown in Fig. 12(a). This step coincides with the SU. If the environment is separable, the cost function is already minimal. In any other case, we can anticipate good starting

¹⁰We also explicitly checked the resulting condition numbers when first a QR and then a LQ, or the other way round, or only a single decomposition is applied. In these cases, one tensor always encounters better condition numbers than the other in the linear equations of the following sweeping.

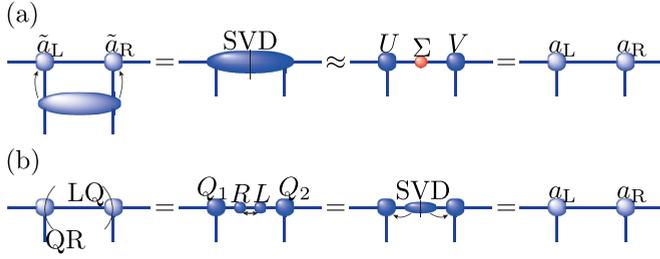


FIG. 12. (Color online) Initial and final steps of the reduced tensor update. (a) Initialization: before we update the reduced tensors by sweeping, we apply a SVD to their joint contraction with the Trotter gate, and keep only the D largest singular values in Σ . Splitting Σ gives the initial tensors $a_L := U\sqrt{\Sigma}$ and $a_R := \sqrt{\Sigma}V$ for the ALS procedure. (b) Final form: after convergence of the ALS sweeping, we put the two tensors on an equal footing.

tensors that are closer to the minimum of the cost function. Second, we optimize the tensors by means of the standard ALS sweeping, in which each tensor update is followed by the standard gauge fixing [16], i.e., the left (right) tensor is QR (LQ) decomposed along its right (left) virtual bond. Third, a gauge choice is made on the internal link of the converged pair as shown in Fig. 12(b).

We have observed that our gauge choices improve the condition number of the norm matrix by several orders of magnitude in all studied cases. This statement is quantified by the results in Table I, which compares typical condition numbers found in the simulation of the Ising and Heisenberg models with and without our gauge fixing. Strictly speaking, the condition number of the norm matrix N_l provides only an upper bound for the final error of the solution \bar{A}_l to the linear system of equations [32]. Therefore a large condition number does not imply low accuracy, but a small condition number implies high accuracy of the solution. In practical computations with finite PEPS, when our gauge transformations are not used, instabilities can occur (e.g., as reported in Ref. [14]) that we have never encountered after our gauge fixing.

TABLE I. We show the mean condition number of the norm matrix with its standard deviation in the reduced tensor update without our gauge fixing, using only the positive approximant, and with our gauge fixing during the FU imaginary time evolution of $D = 2$ (a) and $D = 4$ (b) PEPS of size $N = 11 \times 11$ for the Ising model and of size $N = 10 \times 10$ for the Heisenberg model. The values were obtained averaging over 10 time steps and all tensors in the lattice.

(a)	Model	Positive approximant	Gauge fixing
	$B = 1.0$ Ising	$(2 \pm 3) \times 10^7$	1.1 ± 0.1
	$B = 3.0$ Ising	$(2 \pm 3) \times 10^3$	1.6 ± 0.1
	Heisenberg	$(8 \pm 5) \times 10$	1.08 ± 0.02
(b)	Model	Positive approximant	Gauge fixing
	$B = 1.0$ Ising	$(9 \pm 205) \times 10^{13}$	$(1 \pm 3) \times 10^4$
	$B = 3.0$ Ising	$(4 \pm 158) \times 10^{13}$	$(5 \pm 6) \times 10^2$
	Heisenberg	$(3 \pm 2) \times 10^4$	5 ± 3

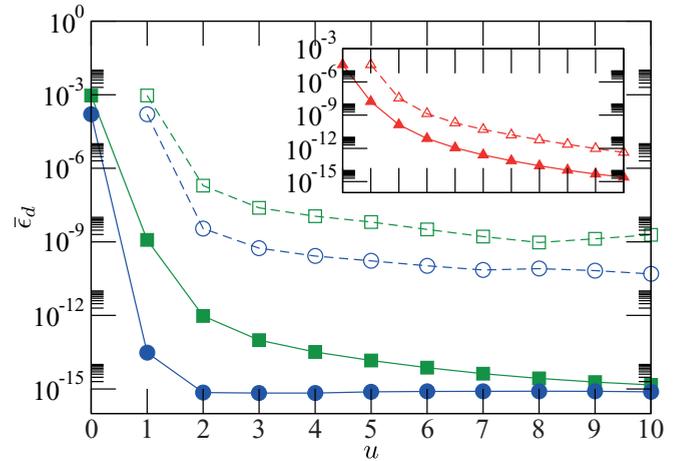


FIG. 13. (Color online) Mean value of the relative change $\epsilon_d(u) := |d(u) - d(u-1)|/|d_{\text{init}}|$ of the cost function d , Eq. (1), after consecutive update sweeps u over a tensor pair computed with respect to the initial value of the cost function d_{init} , for the $D = 4$ reduced tensor update setting of Table I. We compare the FU evolution without our gauge fixing using only the positive approximant (open symbols) to the same propagation with our gauge fixing (filled symbols), for a 11×11 Ising model at $B = 1.0$ (circles) and 3.0 (squares), and for a 10×10 Heisenberg model (triangles).

We can also investigate the effect of our gauge fixing on the convergence of the ALS sweeping, which can be gathered from Fig. 13 for the update of the reduced tensor. Most remarkably, in the presence of the gauge transformations, already the initial SVD drastically reduces the cost function Eq. (1), by a value that in all considered cases is larger than the one attained after one sweep without the gauge transformations. Furthermore, the final total reduction of the cost function is also larger with our gauge fixing than without. Because the relative change of the cost function in Fig. 13 always decreases faster in the presence of our gauge transformations, we conclude that the latter accelerate the convergence of the ALS scheme. Our results indicate that a simplified tensor update consisting of the combination gauge fixing and SVD only, without the ALS sweeps, might be successful. Indeed, for the Ising model, the cost function after our gauge fixing and SVD is already smaller than after ten sweeps without our gauge fixing. However, the sweeping can further decrease the cost function, and this is revealed most evidently for the Heisenberg model.

So far, we assumed that the update is performed on two directly neighboring tensors, after applying on them one of the Trotter gates of a nearest-neighbor Hamiltonian. The discussion can be extended to the update of more distant tensors, as would appear in the case of Hamiltonians with long-range interactions. However, for n nonadjacent reduced tensors the dimension of the norm matrix is typically $(d^2 D^2)^n \times (d^2 D^2)^n$, and its diagonalization, even for two tensors, is not desirable. A further simplification is to choose the gauge transformations from the local environment of a single tensor, since the norm matrix in this case has size $d^2 D^2 \times d^2 D^2$. By means of a numerical simulation we confirmed that such a local gauge choice can produce condition numbers comparable to the ones

obtained from the gauging of the environment of the pair.¹¹ Moreover, we found that the gauge matrices L and R computed from the local norm tensors of each separate tensor in the pair can also be applied to the pair environment, and then we can follow the update procedure of Fig. 12. Thus we expect that local gauge choices similarly improve the tensor update in the case of general long-range interactions, where the tensor initialization and final form in Fig. 12 will be given by their analogues from TEBD [23].

While the discussion here is focused on the reduced tensors, in Appendix B we derive an efficient gauge fixing for the full tensors. This gauge fixing equally improves the condition number of the norm matrix and the convergence of the ALS sweeping in the tensor update of the full tensors. Because all our gauge transformations are derived from and applied to the norm TN alone, they do not explicitly depend on the operator whose action on the PEPS is approximated. We would therefore expect that our gauge choices similarly improve the original time evolution algorithm [1,16] in which the action of projected entangled pair operators (PEPO) on PEPS is approximated.¹²

3. Stability issues

The previously described gauge choices guarantee a better conditioned norm matrix. For the stability, precision, and efficiency of the algorithms, especially when the environment approximation is very rough (e.g., by using small clusters or boundary bond dimensions), also the following factors need to be taken into account.

(1) For PEPS, the matrices N_l are not exactly Hermitian and positive semidefinite. The advisable strategy is to replace them by their closest Hermitian approximants $(N_l + N_l^\dagger)/2$ and additionally set to zero any negative eigenvalues in order to get the closest positive semidefinite approximant of N_l , as described above for the environment of the reduced tensor pair.¹³

(2) In general, some eigenvalues of N_l are zero and its positive subspace is ill-conditioned. That is why N_l^{-1} must be a pseudoinverse. A cutoff is set such that only the subspace of N_l with eigenvalues larger than a certain value is considered in the construction of the pseudoinverse.

¹¹We propagated, using the FU, several 11×11 $D = 2$ PEPS, obtained from the SU for the Ising model with different fields, and we monitored eigenvalues and singular values of norm and gauge matrices throughout the evolution.

¹²We can also expect that the direct variational minimization of the energy benefits from our gauge transformations because they improve the condition number of the norm matrix N_l that enters the generalized eigenproblem $H_l \hat{A}_l = \lambda N_l \hat{A}_l$, which has to be solved for the tensor update at lattice site l .

¹³In the direct variational minimization of the energy, also the matrices H_l are not exactly Hermitian and hence should be replaced by their closest Hermitian approximants $(H_l + H_l^\dagger)/2$. Furthermore, the eigendecomposition of $(N_l + N_l^\dagger)/2$ enables us to replace the generalized eigenproblem by a standard one, $\sqrt{N_l}^{-1} H_l \sqrt{N_l}^{-1} \tilde{B}_l = \lambda \tilde{B}_l$. Its lowest eigenvector, \tilde{B}_l , yields the desired new variational parameters via $\hat{A}_l = \sqrt{N_l}^{-1} \tilde{B}_l$.

(3) Finally, the correct tensor normalization has a decisive impact. Imaginary time evolution steadily modifies the norm of the state. Thus we impose the normalization of the PEPS, $\langle \psi | \psi \rangle = 1$, after each set of Trotter gates, and, in order to avoid the existence of very small or very large tensors, we additionally scale all PEPS tensors to have the same largest element absolute.

IV. PERFORMANCE OF FINITE PEPS

With the aim of analyzing its performance in terms of system size and bond dimension, we have applied the generic finite PEPS code to the ground-state search for the Heisenberg and quantum Ising model, and compared the results to those obtained by other numerical methods, when available. Our best PEPS results were obtained with the FU, i.e., updating the reduced tensors, applying the Trotter gates sequentially, and approximating the full contraction of the environment by means of general boundary MPO. This combination of techniques allowed us to push the simulations to lattices of size up to 21×21 . On the same systems, we ran also the SU for finite PEPS.

A. Convergence procedure

In each case, the PEPS ground-state approximation was found by means of imaginary time evolution. The initial state was always a $D = 2$ PEPS, which was constructed by embedding in it a separable PEPS and replacing the zero entries by small random numbers. Beginning with the time step $\tau = 0.01$, the propagation was performed long enough for the energy to converge, and then the procedure was repeated for smaller time step(s). After convergence was attained for the minimum time step, the scheme was iterated for a larger bond dimension, starting from a previously converged PEPS as initial state.

We observed that the converged SU PEPS of a certain bond dimension was always a good initial state for further propagation with the FU for this bond dimension. On the one hand, in general, the SU PEPS can already be a good ground-state approximation and then only few further steps with the FU are required. On the other hand, we have found, that such state required smaller values of D' when the evolution was continued with the FU.

Energies and correlators reported here for a certain value of D correspond to the final PEPS for the smallest time step. The error in the corresponding observable was estimated via the difference to the expectation value calculated with the converged PEPS for the previous time step. All contractions were performed with boundary bond dimension $D' = 100$, big enough to neglect contraction errors, as we explicitly checked by comparison to results from $D' = 200$.

B. Heisenberg model

We considered imaginary time evolution with an antiferromagnetic Heisenberg Hamiltonian $\hat{H} = \sum_{\langle l,m \rangle} \vec{S}_l \cdot \vec{S}_m$. This model on a two-dimensional square lattice is a paradigmatic benchmark Hamiltonian because quantum Monte Carlo methods provide quasiexact results for very large system sizes [33], and thus we can directly compare our results to quantum Monte

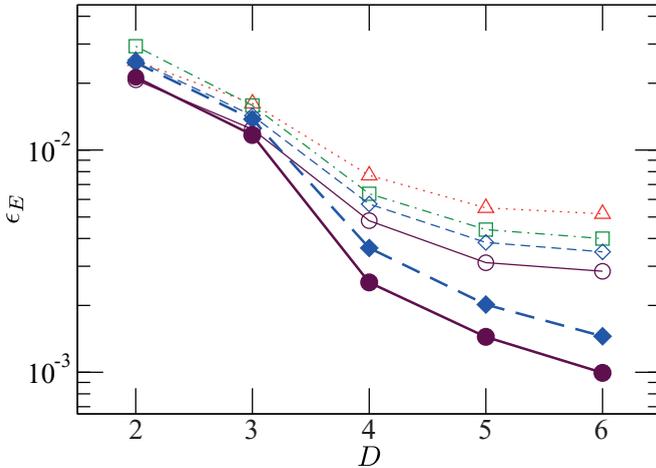


FIG. 14. (Color online) Relative energy error $\epsilon_E := |E(D) - E_0|/|E_0|$, where E_0 denotes the exact ground-state energy from the ALPS library [34–36], of the SU (open symbols) and the FU (filled symbols) for different lattice sizes. In the case of the SU, we consider $N = 10 \times 10$ (triangles), 14×14 (squares), 16×16 (diamonds), and 20×20 (circles). In the case of the FU, we consider $N = 10 \times 10$ (diamonds) and 14×14 (circles).

Carlo simulations.¹⁴ In the context of PEPS, the ground-state order parameter of this model, i.e., the squared staggered magnetization $M_{\text{stag}}^2 := \frac{1}{N^2} \sum_{l,m=1}^N (-1)^{l+m} \langle \vec{S}_l \cdot \vec{S}_m \rangle$, is particularly challenging [37] (also on a honeycomb lattice [11]) and a precise determination has so far only been possible with very large bond dimension $D = 16$ in Ref. [13]. Here we want to find out what our improved algorithmic procedures can do.

To our Heisenberg Hamiltonian we added a small staggered magnetic field $B_Z \sum_l (-1)^l S_l^z$, which we slowly switched off during the evolution, starting from $B_Z = 10^{-3}$. In the presence of this staggered field, the SU(2) symmetry of the Heisenberg model is explicitly broken and smaller values of D' suffice. This procedure improved the convergence of all our algorithms significantly. In the case of the SU, it helped to avoid local minima and reach lower final energies, in particular on the largest 20×20 lattice. And, in the case of the FU, already when the staggered field was still switched on, low values of the energy were attained while smaller values of D' were required. All propagations were performed for time steps $\tau = 10^{-2}$ and 10^{-3} .

Figure 14 shows the convergence of the energy with increasing bond dimension. We observe that, while the FU energy error decreases rapidly with D , the SU energies saturate, and for bond dimensions up to $D = 6$, the lowest SU energies lie between the values for $D = 3$ and 4 obtained with the FU. This is consistent with our earlier observations in Ref. [21] based on smaller lattices. Both the SU and the FU produce better energies when the lattice size increases.

We can now compare our energy accuracies to the existing literature. The original finite PEPS algorithm [1,16] obtained a lowest energy per site -0.62515 on a 10×10 lattice,

using a time step $\tau = 0.001$ and bond dimension $D = 4$ (this PEPS result is given in Refs. [38,39]). For this system size and the same values of τ and D , we now achieve the slightly lower energy per site $-0.62637(2)$, and we can also provide the converged $D = 6$ result $-0.62774(1)$. All our energies as well as our quantum Monte Carlo reference values corresponding to Fig. 14 are collected in Appendix C. Our $D = 4$ energy per site is already lower than the best values reported for the wave function ansatz (block) sequentially generated states (-0.61713) [38], entangled-plaquette states ($-0.6258(1)$) [39], and string bond states (-0.6225) [40], to which we can directly compare because they also considered finite systems with open boundary conditions. For infinite systems, the iPEPS ansatz attains slightly better energy precisions between 10^{-3} and 10^{-4} for $D = 4$ to 6, as reported in Ref. [41]. For large finite cylinders, the best DMRG results are also more accurate: Ref. [42] analyzes the Heisenberg model on a cylinder with a constant staggered magnetic field on the boundaries and, by making use of S^Z symmetry in the algorithm, reaches an energy accuracy of 10^{-4} on a 20×10 lattice.

In order to check the accuracy of the ground-state approximation, we evaluated also nonlocal observables. In particular, we computed the correlator $\langle \vec{S}_l \cdot \vec{S}_{l+x} \rangle$, in the center of the lattice for two sites separated by a distance x , either along the diagonal or along the same column. We checked explicitly that the correlators of the converged PEPS along the diagonal and vertical direction are quantitatively very similar. This feature is obviously due to the PEPS ansatz and would be harder to reproduce, e.g., with MPS in two dimensions. The precision of our considered spin-spin correlator $\langle \vec{S}_l \cdot \vec{S}_{l+x} \rangle$ also indicates the precision that can be expected for the order parameter $M_{\text{stag}}^2 := \frac{1}{N^2} \sum_{l,m=1}^N (-1)^{l+m} \langle \vec{S}_l \cdot \vec{S}_m \rangle$. Since the former quantity, being dependent on the distance x , provides more information than the latter quantity, being just a single number, we focus here on the spin-spin correlator.

The results for the diagonal correlators in 10×10 PEPS are shown in Fig. 15(a), and Fig. 15(b) displays the vertical correlators for 14×14 PEPS. We observe that the FU converges quickly to the true correlator with increasing bond dimension. Although for fixed D the error grows with the distance x , for fixed x it decreases fast with D . In particular, if we consider the correlator at distance $x = L/2$, as commonly done for the construction of the thermodynamic value via finite size scaling, we read off $\epsilon_C^{D=6} \approx 0.01$ and $\epsilon_C^{D=7} \approx 0.003$ on the 10×10 lattice, and we find $\epsilon_C^{D=5} \approx 0.07$ and $\epsilon_C^{D=6} \approx 0.01$ on the 14×14 lattice. As for the energy, the SU results saturate, and they get better when the system size is larger.

We want to compare our results for the spin-spin correlator to previous works. The widely used iPEPS algorithms achieve a remarkably low relative energy error in the thermodynamic limit [41] while their relative correlator error ≈ 0.1 reported in Ref. [37] for $D = 5$ is still rather high (although larger values of D are accessible within iPEPS algorithms nowadays [5,43] by making use of symmetries [44,45]). In Ref. [13], the SU was used together with Monte Carlo sampling to reach much larger bond dimensions, and their best accuracies obtained with $D = 16$ were $0.003(2)$ on a 8×8 lattice and $0.013(2)$ on a 16×16 lattice, assuming periodic boundary conditions. We

¹⁴We obtain our quantum Monte Carlo reference values from the ALPS library [34–36].

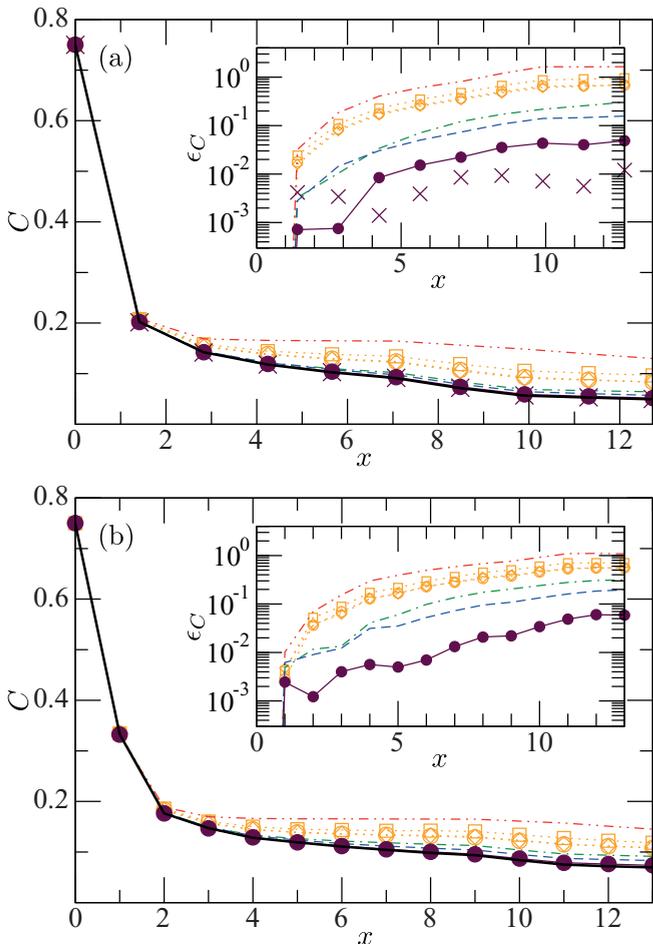


FIG. 15. (Color online) Spin correlations $C(x) := |\langle \vec{S}_i \cdot \vec{S}_{i+x} \rangle|$ (main) and relative error $\epsilon_C(x) := |C(x) - C_0(x)|/|C_0(x)|$ (inset), with the exact values $C_0(x)$ (thick line) from the ALPS library [34–36], for two sites separated by distance x along the diagonal in the center of 10×10 PEPS (a) and along the vertical in the center of 14×14 PEPS (b). We consider PEPS Heisenberg ground-state approximations from the FU with $D = 2$ (dash-double-dotted), 4 (dash-dotted), 5 (dashed), 6 (filled circles), and 7 (crosses), and from the SU with $D = 4$ (squares), 6 (diamonds), and 8 (open circles).

now attain the same precisions here on 10×10 and 14×14 lattices already with much smaller bond dimensions $D = 6$ and 7. Again, the best DMRG results are still more accurate: Ref. [42] reports an uncertainty of 0.0007 for the observable $|\langle S^Z \rangle|$ in the center of a 20×10 cylinder with constant staggered magnetic fields on the boundaries.

We can try to understand the characteristics of the SU and the FU results with the help of the environment approximation used in their tensor updates. As we have argued in Ref. [21], SU and FU represent special cases of a unifying CU_δ : the SU is equivalent to clusters of size $\delta = 0$ in the tensor update, while the FU corresponds to the largest possible cluster size $\delta = L - 1$. Here we showed that the cluster contraction error as a function of the cluster size behaves like the correlation function of the considered PEPS, such that states with short correlation lengths can be accurately contracted by means of small clusters. It is then reasonable to expect that the cluster

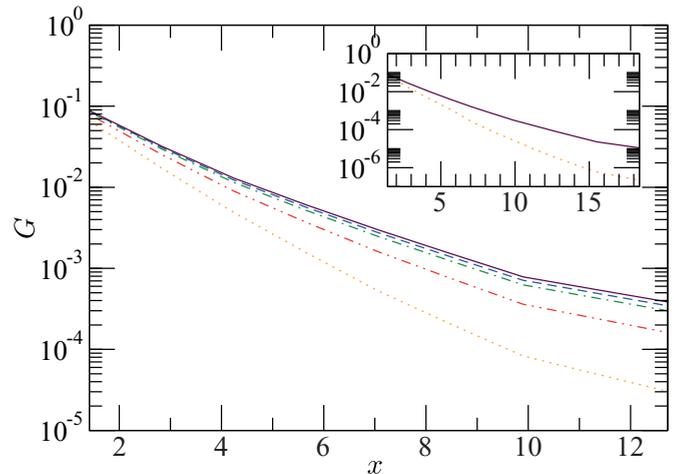


FIG. 16. (Color online) Correlation function $G(x) := \langle \vec{S}_i \cdot \vec{S}_{i+x} \rangle - \langle \vec{S}_i \rangle \cdot \langle \vec{S}_{i+x} \rangle$ for two sites separated by distance x along the diagonal in the center of 10×10 (main) and 14×14 (inset) PEPS Heisenberg ground-state approximations of bond dimension $D = 4$. We compare SU (dotted), CU_1 (dash-double-dotted), CU_2 (dash-dotted), CU_3 (dashed), and FU (solid).

size δ used in CU_δ limits the finally achievable correlation length. We address this question on a 10×10 lattice with $D = 4$ in the main part of Fig. 16. Indeed, the correlation function decays slower when larger clusters are used in the CU.

Moreover, we can gather from Fig. 16 that the correlation functions for system size 14×14 from SU as well as FU decay faster than the corresponding ones for system size 10×10 , while Fig. 14 shows that a higher energy accuracy is attained on the larger lattice. This indicates that, for the finite systems with open boundary conditions considered here, the true correlation length of the Heisenberg model slightly decreases with growing lattice size. In the context of the SU, this would explain why the SU results of Figs. 14 and 15 are better on larger lattices. And in the context of the FU, this would explain our numerical observation that the convergence of energies and spin-spin correlators required smaller values of D' for larger systems:¹⁵ a smaller correlation length can be captured with a smaller cluster size δ in the CU_δ and the contraction precision achieved with such δ can equally be obtained by the full contraction, used in the FU, with correspondingly smaller value of D' (see Figs. 12 and 13 in Ref. [21]).

C. Quantum Ising model

We have also applied our finite PEPS algorithms to the quantum Ising model with transverse field,

¹⁵To be precise, our FU energies and spin-spin correlators for the Heisenberg model were converged with $D' = 2D^2$ for all $D = 2$ to 6 on the 14×14 lattice, i.e., our results did not change anymore when we further ran the FU with larger D' . However, this convergence occurred only with $D' = 75$ for $D = 5$ and $D' = 126$ for $D = 6$ on the 10×10 lattice.

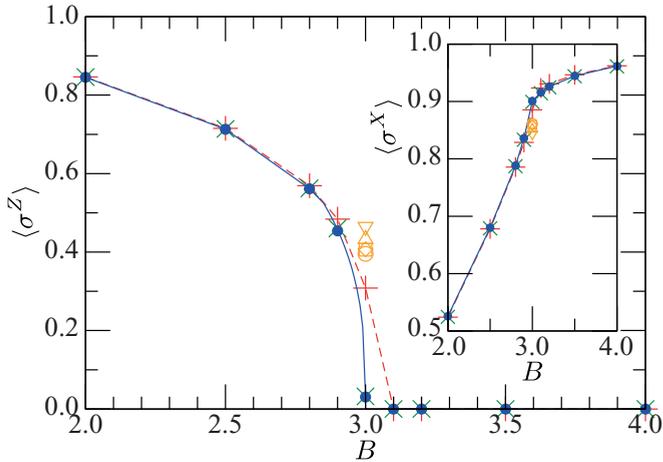


FIG. 17. (Color online) Observables $\langle \sigma^Z \rangle$ (main) and $\langle \sigma^X \rangle$ (inset) evaluated in the center of 21×21 PEPS Ising ground-state approximations from the FU with $D = 2$ (pluses), 3 (crosses), and 4 (filled circles), where the $D = 4$ are basically on top of the $D = 3$ results. Open symbols show the SU at $B = 3.0$, for $D = 2$ (down-triangles), 4 (up-triangles), 5 (squares), 6 (diamonds), and 7 (open circles). We interpolate the FU $D = 4$ $\langle \sigma^Z \rangle$ results between $B = 2.85$ and $B_0 := 3.000035$ with $|B - B_0|^{0.34}$.

$\hat{H} = - \sum_{\langle l,m \rangle} \sigma_l^Z \sigma_m^Z - B \sum_l \sigma_l^X$. This Hamiltonian features a quantum phase transition in the thermodynamic limit, and its critical point $B_c \approx 3.044$ and exponent $\beta \approx 0.327$ are known very accurately thanks to finite size scaling with the quantum Monte Carlo method [46]. Since iPEPS have already very successfully demonstrated the adequacy of the PEPS ansatz for the quantum Ising model even at criticality [9,10] (Ref. [10] reports $B_c \approx 3.04$ and $\beta \approx 0.328$), we present our results here and in Appendix C just for benchmark purposes, e.g., to enable a comparison with another PEPS implementation or with another wave function ansatz. We thus consider here only few different values of the magnetic field around $B = 3$ and run our computations only for the two system sizes 11×11 and 21×21 . For each value of B , we converge the imaginary time evolution independently, using time steps $\tau = 10^{-2}$, 10^{-3} , and 10^{-4} .

Figure 17 shows the order parameter evaluated in the center of 21×21 PEPS from the FU for several points in the phase diagram. Without performing a finite size scaling, we can already extract estimates of the critical point $B_c \approx 3.0$ and exponent $\beta \approx 0.34$ from this finite system, which are close to the iPEPS results [9,10]. We conclude that this lattice is already large enough to display features similar to iPEPS.

For comparison, we also present SU results at $B = 3$. As expected from our previous analysis, the SU does not work well there, where the correlation length should be large. Figure 18 shows that the FU can indeed generate PEPS with larger correlation lengths. While a least squares fit gives a correlation length for the FU $D = 4$ PEPS $\zeta_{\text{FU}}^{D=4} \approx 2.6$, it reveals for the SU $D = 7$ PEPS only $\zeta_{\text{SU}}^{D=7} \approx 1.2$. The inset of Fig. 18 demonstrates the largest correlation length $\zeta_{\text{FU}}^{D=4} \approx 4.3$ for the 11×11 lattice. Notice that, here, we have not performed such an extensive convergence analysis with D' as we have

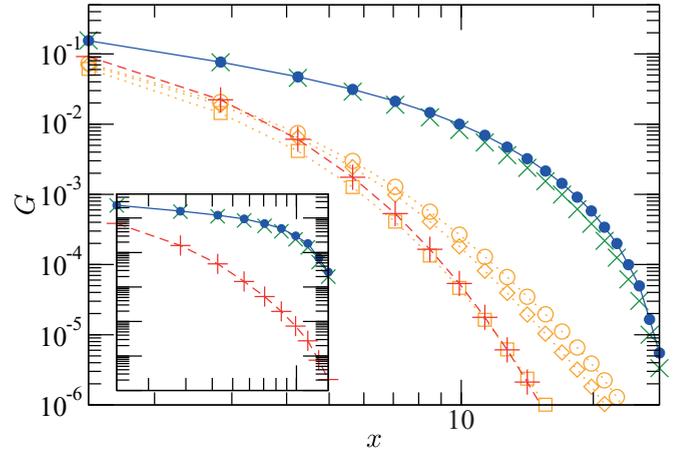


FIG. 18. (Color online) Correlation function $G(x) := \langle \sigma_l^Z \sigma_{l+x}^Z \rangle - \langle \sigma_l^Z \rangle \langle \sigma_{l+x}^Z \rangle$ for two sites separated by distance x along the diagonal in the center of PEPS ground-state approximations of the Ising model on a 21×21 lattice with transverse field $B = 3.0$ (main) and on a 11×11 lattice with $B = 2.8$ (inset). The results were obtained with the FU using bond dimension $D = 2$ (pluses), 3 (crosses), and 4 (filled circles), and with the SU using $D = 4$ (squares), 6 (diamonds), and 7 (open circles).

done before for the Heisenberg Hamiltonian.¹⁶ Nevertheless, we want to emphasize that our correlation functions for the 21×21 lattice are in perfect agreement with the best iPEPS results [10].

Remarkably, long correlation lengths can be analyzed, i.e., large clusters can be contracted, with very high accuracy in the framework of PEPS. This constitutes clear evidence for the power of general boundary MPO. They can capture the correlations of a large cluster size δ with a boundary bond dimension much smaller than the one needed for the exact contraction, $D' = D^{2\delta}$.

V. CONCLUSIONS

In this paper, we have reviewed various aspects that need to be taken into account in the implementation of efficient state-of-the-art finite PEPS algorithms. Within the two main parts of PEPS algorithms, namely the environment approximation and the tensor update, we have analyzed algorithmic strategies that improve the efficiency and stability of the procedures, and the physical properties of the solution.

The environment approximation has decisive influence on the precision of the final PEPS of an imaginary time evolution, and is equally crucial for the computation of expectation values. We have shown how the accuracy of the cluster strategy [21], which allows for a natural trade-off between precision and computational cost of the environment, is fundamentally connected to the correlation length of the

¹⁶For the quantum Ising model and all considered transverse fields B , on the 11×11 lattice, we ran the FU with $D' = 8$ for $D = 2$, $D' = 54$ for $D = 3$, and $D' = 128$ for $D = 4$, while on the 21×21 lattice, we ran the FU with $D' = 8$ for $D = 2$, $D' = 36$ for $D = 3$, and $D' = 128$ for $D = 4$.

state. Additionally, we have demonstrated that it is possible to make use of purification MPO in order to ensure a positive environment approximation, and that this overcomes the limitations of the single-layer algorithm [14]. The numerical techniques analyzed in this problem can straightforwardly be applied to the cluster update [21] and to the full update, but also to other scenarios where a positive MPO is required, e.g., to describe the mixed state of a one-dimensional system.

Not only the environment approximation, but also the method chosen for the tensor update affects the cost and stability of the routines. We have proposed an update scheme that is more efficient and better conditioned than the one from the original algorithm [1,16]. By restricting the variational parameters to the reduced tensor, the update is drastically accelerated. For both the reduced and the full tensor, we have formulated gauge fixings that significantly improve the conditioning. These gauge fixings, additionally, when combined with a cheap SVD, constitute a promising simplified but fast tensor update procedure.

Finally, we have combined the ingredients discussed above in an efficient implementation of finite PEPS imaginary time evolution, capable of dealing with large systems and bond dimensions. In particular we have opted for the sequential application of Trotter gates, using general boundary MPO in the contraction of the full environment,¹⁷ and restricting the update to the reduced tensors. To benchmark the performance of finite PEPS and to quantitatively assess the algorithmic properties, we have applied the code to the ground-state search for the Heisenberg and the quantum Ising model.

We have presented ground-state calculations for system sizes up to 21×21 and bond dimensions up to $D \approx 7, 8$. Our results demonstrate the adequacy of the PEPS ansatz for the description of strongly correlated quantum many-body systems, with energy and order parameter converging fast with increasing bond dimension, when they were obtained with the full update. In that case, thanks to the algorithmic improvements developed in this article, we have been able to achieve precisions of the spin-spin correlator in the Heisenberg model using bond dimensions $D = 6$ and 7 that previously had only been attained using a much larger $D = 16$ in Ref. [13]. Our analysis of a 21×21 quantum Ising model gave, already without finite size scaling, critical point, critical exponent, and correlation functions in good agreement with the iPEPS results [9,10].

The simple update [8] and the cluster update [21] using small cluster sizes, while ensuring a less costly environment and thus being able to deal with larger bond dimensions, do not produce the best ground-state approximation for a certain value of D , and, in particular, give rise to PEPS with limited correlation length, which is especially relevant for strongly correlated systems as, e.g., the Ising model close to criticality. This makes clear that the largest bond dimension attained is not the significant measure of the power of a PEPS algorithm.

By reaching system sizes typically considered for finite size scaling, we have given evidence that finite PEPS, when

all algorithmic details are taken into account, offer a feasible unbiased alternative to their infinite counterpart iPEPS [9]. On the other hand, the algorithmic methods proposed here can also be applied to iPEPS, and the feasibility of large finite PEPS demonstrated here suggests that large unit cells are possible in iPEPS, such that their potential bias due to a finite unit cell can be well analyzed by systematically increasing the unit cell from small to very large size.

Our analysis has been carried out with a generic implementation of PEPS algorithms, so that one can expect that adapting the methods to the specific properties of a certain problem will further enhance the performance. A particularly promising next step is to incorporate the symmetries of the considered Hamiltonian in the tensors [44,45], a key element of ground-breaking two-dimensional DMRG studies, such as Refs. [47,48], and of seminal iPEPS calculations, such as Refs. [5,43].

ACKNOWLEDGMENTS

M. L. thanks B. Bauer, P. Corboz, G. De las Cuevas, I. Pižorn, and L. Wang for discussions. He wants to particularly thank S. Iblisdir, V. Murg, R. Orús, and M. Rizzi for providing key insights into algorithmic aspects. This project was partly funded by the **EU through SIQS** grant (FP7 600645) and the DFG (NIM cluster of excellence). All authors thank the Pedro Pascual Benasque Center for Science (CCBPP), where part of this work was carried out.

APPENDIX A: PURIFICATION APPROXIMATIONS

Approximating the boundary by a purification MPO, as described in Sec. III A 2, requires the solution of nonlinear equations for each tensor A_l . Different algorithms can be used for this purpose, and we have tried and compared three methods.

(1) Linearization. Instead of solving the equations for the product $A_l A_l^*$, we solve them for $A_l B_l$, treating A_l and A_l^* as independent tensors. In order to achieve convergence, the change of the tensor in each iteration needs to be small, and hence we construct the solution of the i th iteration according to $A_l^{(i)} = (1 - \alpha)A_l^{(i-1)} + \alpha A_l$, where A_l solves the linearized equations of the previous iteration and is added to the previous solution $A_l^{(i-1)}$ with a weight α . The latter parameter must be chosen small enough to guarantee a decreasing cost function, and large enough to avoid unnecessarily long convergence times.¹⁸ The construction of the individual parts of the linear equations has the leading cost $\mathcal{O}(dD^6 D'^4) + \mathcal{O}(D^4 D'^6)$. Because we cannot impose a kind of canonical form that gives a trivial norm matrix $N_l = 1$, we have to explicitly contract its tensor network, which contributes a cost $\mathcal{O}(d'^2 D^2 D'^6)$. Finally, computing the pseudoinverse to solve the linear equations requires $\mathcal{O}(d'^3 D^3 D'^6)$ operations, which typically represent the dominant cost when $d' \geq D$.

¹⁷Notice that clusters could be beneficial in a parallel implementation as discussed in Ref. [21].

¹⁸We found that the optimization worked well if the initial value was $\alpha = 0.01$ and was multiplied by 0.8 whenever the cost function increased.

TABLE II. We show the mean condition number of the norm matrix with its standard deviation in the full tensor update, for $D = 2$ and the setting of Table I.

Model	Positive approximant	Gauge fixing
$B = 1.0$ Ising	$(2 \pm 3) \times 10^9$	1.3 ± 0.2
$B = 3.0$ Ising	$(2 \pm 2) \times 10^4$	2.9 ± 0.7
Heisenberg	$(1.3 \pm 0.8) \times 10^3$	1.15 ± 0.05

(2) Conjugate gradient. We employ a canned routine¹⁹ that comprises a conjugate gradient method with line minimization. It has the lowest computational cost, as it only requires the computation of the cost function and its gradient with respect to a single tensor, which can be obtained with $\mathcal{O}(dD^6 D''^4) + \mathcal{O}(D^4 D''^6)$ operations.

(3) Newton method. It approaches a root of the gradient by iterating $\mathcal{H}_l^{(i-1)}(\vec{A}_l^{(i)} - \vec{A}_l^{(i-1)}) = -\vec{G}_l^{(i-1)}$, where $\mathcal{H}_l^{(i-1)}$ denotes the Hessian matrix and $\vec{G}_l^{(i-1)}$ the gradient of the cost function with respect to the tensor components at site l , evaluated with the solution $A_l^{(i-1)}$ of the previous iteration.²⁰ The Newton method has the advantage that the step width is naturally given, in contrast to the linearized equations where α needs to be chosen heuristically, and in contrast to the conjugate gradient routine where it is determined via line search. In addition to the parts of the conjugate gradient algorithm, the Newton method needs the Hessian matrix of the cost function, which contributes $\mathcal{O}(d^2 D^2 D''^6)$ to the cost, and its pseudoinverse, determined by $\mathcal{O}(d^3 D^3 D''^6)$ operations, such that the leading cost $\mathcal{O}(dD^6 D''^4) + \mathcal{O}(D^4 D''^6) + \mathcal{O}(d^3 D^3 D''^6)$ is the same as for the linearized problem.

To compare the different alternatives, we benchmarked their performance in the search for an optimal purification with fixed $D'' = 2$ and varying d' , given a reference purification with $D'' = 4$ and $d' = 4$. The latter was constructed by taking two rows from one edge of a PEPS norm TN, for several 11×11 $D = 2$ SU ground-state approximations of the Ising model at various magnetic fields. As a general rule, the initial tensors for the search with incremented purification bond $d' + 1$ were chosen as the previous solution for d' where the extra elements were filled with uniformly distributed random numbers. From the three considered algorithms, the Newton method performed best. It converged reliably for all d' within few local updates per tensor.

All the methods benefit from initial variables that are already close to the final solution. A sensible numerical approach to purification approximations can then be implemented in two steps: firstly, the computation of the optimal purification via the SL algorithm [14], and secondly, the further optimization of that purification via the Newton method.

¹⁹For conjugate gradient minimization we use `nag_opt_conj_grad` from the NAG library [49].

²⁰We observed that, occasionally, the Hessian matrix had many negative eigenvalues and then it was crucial to include only the positive ones (above a certain cutoff) in the construction of its pseudoinverse.

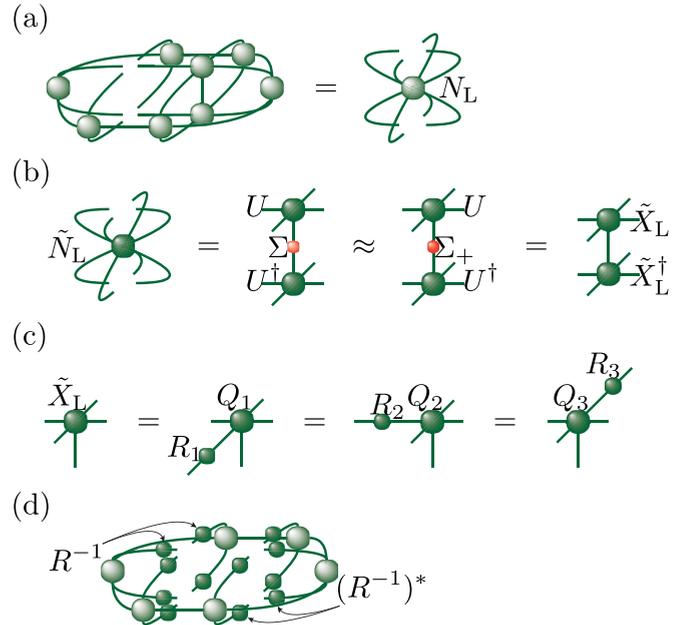


FIG. 19. (Color online) Gauge fixing on the environment tensors of the full tensor pair, when the environment is nonseparable. (a) The environment tensor N_L of the left full tensor. (b) We determine the positive approximant for the environment of the left full tensor via a diagonalization of the Hermitian approximant $\tilde{N}_L := (N_L + N_L^\dagger)/2 = U \Sigma U^\dagger$, in which, then, the negative eigenvalues are discarded in Σ_+ , and, finally, the environment is written as $\tilde{X}_L \tilde{X}_L^\dagger$ in terms of its square root $\tilde{X}_L := U \sqrt{\Sigma_+}$. (c) We perform three independent QR decompositions on \tilde{X}_L . (d) After equally having carried out the previous steps (a) to (c) with the right full tensor, we have six different matrices R . Their inverses are contracted with the corresponding tensors of the boundary MPO.

APPENDIX B: GAUGE FIXING FOR THE FULL TENSORS

When the environment is nonseparable, and the update of the full tensors is considered, gauge transformations can be efficiently computed in such a way that an eigendecomposition of the $D^6 \times D^6$ dimensional norm environment of the pair is not necessary. Instead, the $D^4 \times D^4$ dimensional environments of the left (N_L) and right (N_R) tensor are independently computed [Fig. 19(a)] and replaced by their positive approximants [Fig. 19(b)] like in previous cases. Their square roots are used to obtain the desired gauge transformations for each of the virtual bonds of the pair [Fig. 19(c)]. On each virtual bond we then insert the corresponding product $R^{-1}R$ and absorb the R matrices in the full tensors and their inverses in the environment [Fig. 19(d)].

The update of the full tensor pair proceeds in the way explained in Fig. 20, analogously to the reduced tensor update. As in the latter context, if the environment is separable, the

TABLE III. Energy per site of the Heisenberg model on square lattices of various sizes from quantum Monte Carlo, computed with the ALPS library [34–36].

10×10	14×14	16×16	20×20
$-0.628656(2)$	$-0.639939(2)$	$-0.643531(2)$	$-0.648607(1)$

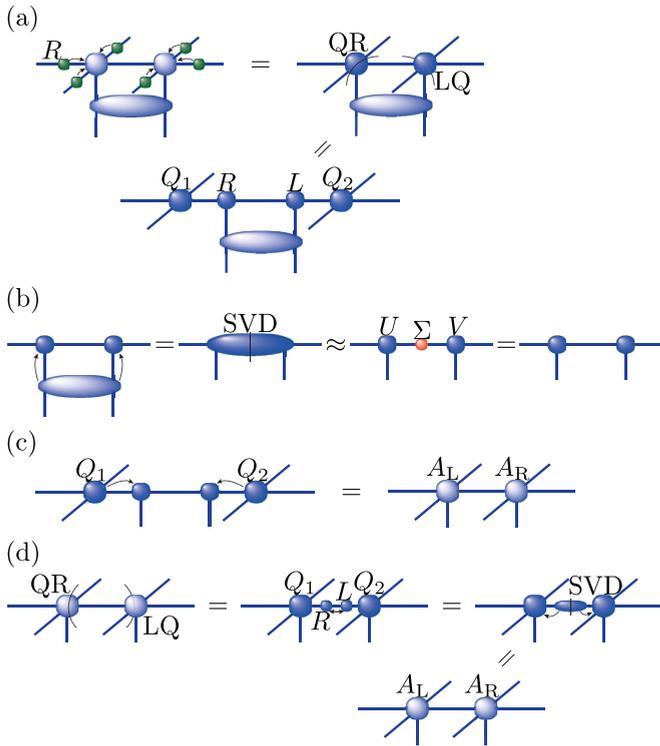


FIG. 20. (Color online) Like in the reduced case of Fig. 12, the update of the full tensor pair also consists of the three stages initialization, optimization, and final form. The optimization is the standard ALS sweeping, in which each full tensor is gauged after its update in the standard way, i.e., the left tensor is QR decomposed along its right virtual bond and the right tensor is LQ decomposed along its left virtual bond. (a) Initialization I. Firstly, we contract the gauge transformations from Fig. 19 with the full tensors, and split off their reduced parts. (b) Initialization II. Secondly, we construct new reduced tensors from a SVD on the tensor pair and the Trotter gate, equally sharing the D largest singular values between the left and right tensor. (c) Initialization III. We recover the full tensors A_L and A_R , which are now the initial tensors for the optimization via ALS sweeping. (d) Final form. After convergence of the sweeping, we put the two tensors on the same footing.

tensor initialization Fig. 20(a) to (c) already minimizes the cost function, while, if the environment is close to separable, we can expect a significant decrease of the cost function. In general, we can anticipate good starting tensors for the following ALS sweeping.

Table II contains typical condition numbers of the norm matrix in the full tensor update without our gauge fixing, using

TABLE IV. Energy per site of PEPS Heisenberg ground-state approximations from the FU.

D	10×10	14×14
2	-0.61310(2)	-0.62631(1)
3	-0.61999(1)	-0.63246(1)
4	-0.62637(2)	-0.63832(3)
5	-0.62739(1)	-0.63901(1)
6	-0.62774(1)	-0.63930(1)

TABLE V. Energy per site of PEPS Heisenberg ground-state approximations from the SU.

D	10×10	14×14	16×16	20×20
2	-0.61281(1)	-0.62115(1)	-0.62719(2)	-0.63519(2)
3	-0.61846(2)	-0.62977(1)	-0.63433(1)	-0.64056(2)
4	-0.62382(1)	-0.63587(1)	-0.63985(1)	-0.64549(2)
5	-0.62520(2)	-0.63713(2)	-0.64106(1)	-0.64659(2)
6	-0.62541(2)	-0.63738(2)	-0.64129(2)	-0.64676(2)

TABLE VI. Energy per site of 11×11 PEPS Ising ground-state approximations from the FU for different transverse fields B .

D	2.0	2.5	2.8
2	-2.40075(1)	-2.74230(2)	-2.98947(5)
3	-2.40076(1)	-2.74243(1)	-2.99094(2)
4	-2.40076(1)	-2.74243(1)	-2.99099(1)
D	2.9	3.0	3.1
2	-3.07945(5)	-3.17128(4)	-3.26400(4)
3	-3.08071(1)	-3.17210(1)	-3.26457(1)
4	-3.08073(1)	-3.17210(1)	-3.26457(1)
D	3.2	3.5	4.0
2	-3.35744(4)	-3.64097(3)	-4.12064(2)
3	-3.35785(1)	-3.64116(1)	-4.12071(1)
4	-3.35785(1)	-3.64116(1)	-4.12071(1)

TABLE VII. Energy per site of 21×21 PEPS Ising ground-state approximations from the FU for different transverse fields B .

D	2.0	2.5	2.8
2	-2.45219(1)	-2.77340(2)	-3.00705(4)
3	-2.45219(1)	-2.77346(1)	-3.00737(1)
4	-2.45219(1)	-2.77346(1)	-3.00737(1)
D	2.9	3.0	3.1
2	-3.09228(5)	-3.18128(6)	-3.27326(4)
3	-3.09287(1)	-3.18242(1)	-3.27406(1)
4	-3.09287(1)	-3.18243(1)	-3.27406(1)
D	3.2	3.5	4.0
2	-3.36617(4)	-3.64849(3)	-4.12685(2)
3	-3.36672(1)	-3.64873(1)	-4.12694(1)
4	-3.36672(1)	-3.64873(1)	-4.12694(1)

TABLE VIII. Energy per site of 21×21 PEPS Ising ground-state approximations from the SU for transverse field $B = 3.0$.

D	3.0
2	-3.1792(4)
3	-3.1806(4)
4	-3.1807(4)
5	-3.1812(5)
6	-3.1812(4)
7	-3.1814(5)

only the positive approximant, and with our gauge fixing. Our gauge fixing improves the condition number drastically.

APPENDIX C: FINITE PEPS ENERGIES

Here, we collect some precise energy values obtained with the PEPS ground-state approximations considered in this

paper. In the case of the Heisenberg model, we compare our results to energies from the quantum Monte Carlo loop algorithm of the ALPS library [34–36], summarized in Table III. The presented values and errors correspond to temperature $T = 10^{-4}$, and they agree with the ones corresponding to $T = 10^{-3}$ within the error bars.

-
- [1] F. Verstraete and J. I. Cirac, [arXiv:cond-mat/0407066](#).
- [2] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [3] U. Schollwöck, *Ann. Phys.* **326**, 96 (2011).
- [4] P. Corboz, S. R. White, G. Vidal, and M. Troyer, *Phys. Rev. B* **84**, 041108(R) (2011).
- [5] P. Corboz, T. M. Rice, and M. Troyer, [arXiv:1402.2859](#).
- [6] Z. Y. Xie, J. Chen, J. F. Yu, X. Kong, B. Normand, and T. Xiang, *Phys. Rev. X* **4**, 011025 (2014).
- [7] M. Levin and C. P. Nave, *Phys. Rev. Lett.* **99**, 120601 (2007).
- [8] H. C. Jiang, Z. Y. Weng, and T. Xiang, *Phys. Rev. Lett.* **101**, 090603 (2008).
- [9] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **101**, 250602 (2008).
- [10] R. Orús and G. Vidal, *Phys. Rev. B* **80**, 094403 (2009).
- [11] Z. Y. Xie, H. C. Jiang, Q. N. Chen, Z. Y. Weng, and T. Xiang, *Phys. Rev. Lett.* **103**, 160601 (2009).
- [12] P. Corboz, R. Orús, B. Bauer, and G. Vidal, *Phys. Rev. B* **81**, 165104 (2010).
- [13] L. Wang, I. Pižorn, and F. Verstraete, *Phys. Rev. B* **83**, 134421 (2011).
- [14] I. Pižorn, L. Wang, and F. Verstraete, *Phys. Rev. A* **83**, 052321 (2011).
- [15] Z. Y. Xie, J. Chen, M. P. Qin, J. W. Zhu, L. P. Yang, and T. Xiang, *Phys. Rev. B* **86**, 045139 (2012).
- [16] F. Verstraete, V. Murg, and J. I. Cirac, *Adv. Phys.* **57**, 143 (2008).
- [17] V. Murg, F. Verstraete, and J. I. Cirac, *Phys. Rev. A* **75**, 033605 (2007).
- [18] V. Murg, F. Verstraete, and J. I. Cirac, *Phys. Rev. B* **79**, 195119 (2009).
- [19] M. Lubasch, V. Murg, U. Schneider, J. I. Cirac, and M.-C. Bañuls, *Phys. Rev. Lett.* **107**, 165301 (2011).
- [20] A. Hosseinkhani, B. G. Dezfouli, F. Ghasemipour, A. T. Rezakhani, and H. Saberi, *Phys. Rev. A* **89**, 062324 (2014).
- [21] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, *New J. Phys.* **16**, 033014 (2014).
- [22] R. Orús, *Ann. Phys.* **349**, 117 (2014).
- [23] G. Vidal, *Phys. Rev. Lett.* **91**, 147902 (2003).
- [24] A. J. Daley, C. Kollath, U. Schollwöck, and G. Vidal, *J. Stat. Mech.* (2004) P04005.
- [25] G. Vidal, *Phys. Rev. Lett.* **93**, 040502 (2004).
- [26] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, *Phys. Rev. Lett.* **93**, 207204 (2004).
- [27] J. I. Cirac, D. Poilblanc, N. Schuch, and F. Verstraete, *Phys. Rev. B* **83**, 245134 (2011).
- [28] G. De las Cuevas, N. Schuch, D. Pérez-García, and J. I. Cirac, *New J. Phys.* **15**, 123021 (2013).
- [29] M. Kliesch, D. Gross, and J. Eisert, [arXiv:1404.4466](#).
- [30] J. B. Keller, *Math. Mag.* **48**, 192 (1975).
- [31] H. Kalis, D. Klagges, R. Orús, and K. P. Schmidt, *Phys. Rev. A* **86**, 022317 (2012).
- [32] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. (The Johns Hopkins University Press, Baltimore and London, 1996).
- [33] A. W. Sandvik, *Phys. Rev. B* **56**, 11678 (1997).
- [34] S. Todo and K. Kato, *Phys. Rev. Lett.* **87**, 047203 (2001).
- [35] A. F. Albuquerque *et al.*, *J. Magn. Magn. Mater.* **310**, 1187 (2007).
- [36] B. Bauer *et al.*, *J. Stat. Mech.* (2011) P05001.
- [37] B. Bauer, G. Vidal, and M. Troyer, *J. Stat. Mech.* (2009) P09006.
- [38] M.-C. Bañuls, D. Pérez-García, M. M. Wolf, F. Verstraete, and J. I. Cirac, *Phys. Rev. A* **77**, 052306 (2008).
- [39] F. Mezzacapo, N. Schuch, M. Boninsegni, and J. I. Cirac, *New J. Phys.* **11**, 083026 (2009).
- [40] A. Sfondrini, J. Cerrillo, N. Schuch, and J. I. Cirac, *Phys. Rev. B* **81**, 214426 (2010).
- [41] E. M. Stoudenmire and S. R. White, *Annu. Rev. Condens. Matter Phys.* **3**, 111 (2012).
- [42] S. R. White and A. L. Chernyshev, *Phys. Rev. Lett.* **99**, 127004 (2007).
- [43] P. Corboz and F. Mila, *Phys. Rev. Lett.* **112**, 147203 (2014).
- [44] S. Singh, R. N. C. Pfeifer, and G. Vidal, *Phys. Rev. B* **83**, 115125 (2011).
- [45] B. Bauer, P. Corboz, R. Orús, and M. Troyer, *Phys. Rev. B* **83**, 125106 (2011).
- [46] H. W. J. Blöte and Y. Deng, *Phys. Rev. E* **66**, 066110 (2002).
- [47] S. Yan, D. A. Huse, and S. R. White, *Science* **332**, 1173 (2011).
- [48] S. Depenbrock, I. P. McCulloch, and U. Schollwöck, *Phys. Rev. Lett.* **109**, 067201 (2012).
- [49] The NAG Library, The Numerical Algorithms Group (NAG), Oxford, United Kingdom (2013), <http://www.nag.com>.