Three-dimensional isometric tensor networks

Maurits S. J. Tepaske^{*} and David J. Luitz[†]

Max Planck Institute for the Physics of Complex Systems, Noethnitzer Strasse 38, 01167 Dresden, Germany

(Received 2 June 2020; revised 3 June 2021; accepted 3 June 2021; published 22 June 2021)

Tensor network states are expected to be good representations of a large class of interesting quantum many-body wave functions. In higher dimensions, their utility is however severely limited by the difficulty of contracting the tensor network, an operation needed to calculate quantum expectation values. Here we introduce a method for the time evolution of three-dimensional isometric tensor networks which respects the isometric structure and therefore renders contraction simple through a special canonical form. Our method involves a tetrahedral site splitting which allows one to move the orthogonality center of an embedded tree tensor network in a simple cubic lattice to any position. Using imaginary time evolution to find an isometric tensor network representation of the ground state of the three-dimensional transverse field Ising model across the entire phase diagram, we perform a systematic benchmark study of this method in comparison with exact Lanczos and quantum Monte Carlo results. We show that the obtained energy matches the exact ground-state result accurately deep in the ferromagnetic and polarized phases, while the regime close to the critical point requires larger bond dimensions. This behavior is in close analogy with the two-dimensional case, which we also discuss for comparison.

DOI: [10.1103/PhysRevResearch.3.023236](https://doi.org/10.1103/PhysRevResearch.3.023236)

I. INTRODUCTION

The Hilbert space dimension of quantum many-body systems grows exponentially with the number of constituents, making the direct handling of many-body wave functions impractical for large systems. Tensor networks are an attempt to tame the many-body wave function, by expressing it in terms of local tensors, which are contracted according to the network structure. This reduces the complexity from an exponential to a polynomial number of variables. While in principle any wave function can be expressed as a tensor network, some particularly entangled states require exponentially large tensors. Fortunately, the manifold of wave functions expressible with small tensor networks includes wave functions with area law entanglement, which are expected to be relevant for the description of ground states of many local quantum many-body systems [1–3].

Tensor network states are particularly successful in one dimension, where they are known as “matrix-product states” (MPS) [4], which have become state-of-the-art machinery for the classical simulation of one-dimensional (1D) many-body systems. This popularity rests primarily on the existence of powerful algorithms to variationally optimize the energy of the state [e.g., the density matrix renormalization group (DMRG) [5]] and on the ability to compute ma-

trix elements of local operators $\langle \phi | \hat{O} | \psi \rangle$ both exactly and efficiently. In particular the second property does not generalize to the higher-dimensional variants of MPS known as “projected-entangled pair states” (PEPS) [6]. It turns out that the exact calculation of a local correlator in an arbitrary PEPS state $\langle \text{PEPS} | \hat{O} | \text{PEPS} \rangle$ —requiring the contraction of a higher-dimensional network—is generally inefficient for generic finite PEPS with open boundary conditions (OBC) already in two dimensions [6–10]. While PEPS are readily formulated in three dimensions (see Fig. 1), currently no efficient contraction method is known. So even though PEPS are efficient representations of area-law entangled quantum many-body wave functions, it is often difficult to extract useful information from them.

The central problem for the generalization of powerful 1D methods to higher dimensions is caused by the fact that cutting a bond in a higher-dimensional PEPS does not separate the network into two disconnected pieces, in contrast to 1D MPS. In MPS methods, the separation of the network into unique “left” and “right” parts by cutting any bond is exploited by using an orthonormal basis to represent the left/right states, and one can then decimate the basis to the dominant components by truncating to the largest singular values in an optimal way [11]. This property is the foundation of MPS evolution algorithms [12]. The absence of such separability in higher-dimensional PEPS diminishes the effectiveness of purely local evolution algorithms, where in the case of a nearest-neighbor interacting system the tensor network is optimized by iterating over the bonds and applying a two-body gate to each bond followed by a truncation of this bond according to the standard time-evolving block decimation (TEBD) [4,11]. Instead, optimal truncation and hence optimal evolution require each gate to be accompanied by a contraction of the full network (dubbed “full update”), which is inefficient as it

^{*}mtepaske@pks.mpg.de

[†]dluitz@pks.mpg.de

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

FIG. 1. A generic 3D PEPS ansatz for the cubic lattice, where the tensors $T_{i\alpha\beta\gamma\delta\kappa\lambda}^{\sigma_i}$ are represented by spheres. The blue legs denote the physical degrees of freedom σ_i and the gray legs denote the virtual degrees of freedom. The connections depict contractions between the virtual legs of neighboring tensors.

generally scales exponentially with network size when performed exactly [8].

One way around the inefficiency of full contraction is to instead perform the contraction approximately [6,8–10], sacrificing precision for speed. Recently there appeared multiple works [13–16] which suggest an attractive alternative: to construct finite PEPS in an *explicit canonical form* in which it can be contracted both exactly and efficiently in a way that local truncation again becomes optimal just like for MPS, thereby circumventing the mentioned problems that occur when dealing with generic PEPS. This does induce a loss of generality, restricting its subspace in Hilbert space to a subspace of generic PEPS, thereby reducing the expressivity of the network [17]. While the effect of this restriction is not yet clear, it becomes irrelevant in the limit of large bond dimensions and therefore seems at least in principle controllable.

In [13] a class of finite two-dimensional (2D) PEPS called “isometric tensor network states” (isoTNS) was introduced for which $\langle \text{PEPS} | \text{PEPS} \rangle$ reduces to a canonical MPS norm, and which can be time evolved using an efficient local evolution algorithm called TEBD^2 . Here we will generalize the isoTNS ansatz to three dimensions and develop an extension of TEBD^2 which we call TEBD^3 . This upgrade to a higher spatial dimension is an important step in developing efficient techniques to simulate generic three-dimensional (3D) quantum many-body systems, especially for cases which are not accessible to quantum Monte Carlo methods due to a sign problem. This importance is stressed by the limited number of existing finite 3D PEPS algorithms [18,19] and generic simulation methods for 3D quantum many-body systems in general [20–22].

II. METHOD

A generic finite 3D PEPS ansatz for a 3D many-body spin-1/2 system can be written in the local basis $\sigma_i = \pm 1$ as

$$|\text{PEPS}\rangle = \sum_{\sigma_1 \dots \sigma_N} \mathcal{C}(T_1^{\sigma_1} \dots T_N^{\sigma_N}) |\sigma_1 \dots \sigma_N\rangle, \quad (1)$$

where $T_i^{\sigma_i}$ represents the set of tensors which contain the complex-valued variational parameters and which are spatially arranged like the spins σ_i . Here $\mathcal{C}$ indicates that all tensors are contracted, giving complex scalar coefficients, which is usually done by choosing the amount of virtual degrees of freedom per $T_i^{\sigma_i}$ equal to the lattice connectivity and then contracting nearest neighbors. In Fig. 1 this is illustrated for a cubic lattice with open boundary conditions, where the pairs of virtual degrees of freedom are represented by the gray bonds and the physical (spin) degrees of freedom are represented by the blue free legs, i.e., in a particular basis we get the tensors $T_{i\alpha\beta\gamma\delta\kappa\lambda}^{\sigma_i}$.

In order to calculate $\langle \text{PEPS} | \text{PEPS} \rangle$ we would contract this tensor network with the physical legs of its conjugate, which we would then have to contract down to a scalar.

A. General properties of isoTNS

The goal of this paper is to design a type of 3D tensor network that allows for the full network contraction to be done exactly and efficiently. To this end we impose an additional structure on the PEPS shown in Fig. 1, such that $\langle \text{PEPS} | \text{PEPS} \rangle = \langle \text{MPS} | \text{MPS} \rangle$ becomes manifest. In particular, we choose the majority of the T_i to be isometric, meaning that these T_i reduce to identities when contracted with their conjugate over a subset of the legs, e.g.,

$$\sum_{i\alpha\beta\gamma\delta\kappa} T_{\sigma_i}^{\dagger\alpha\beta\gamma\delta\kappa\lambda} T_{\sigma_i}^{\alpha\beta\gamma\delta\kappa\eta} = \mathbb{1}^{\lambda\eta} \quad (2)$$

and

$$\sum_{i\alpha\beta\gamma\delta} T_{\sigma_i}^{\dagger\alpha\beta\gamma\delta\kappa\lambda} T_{\sigma_i}^{\alpha\beta\gamma\delta\nu\eta} = \mathbb{1}^{\kappa\nu} \mathbb{1}^{\lambda\eta}. \quad (3)$$

If T is unitary instead of isometric, these constraints also hold under $T \rightarrow T^\dagger$, which in the case of an isometry instead gives a projector.

In the language of tensor network diagrammatics we can represent these constraints by decorating the legs with arrows [4,13], where incoming arrows represent contracted indices in the isometry constraint and where outgoing arrows represent free indices. In Fig. 2 we illustrate this notation for tensors with four virtual legs. Here the diagrams in the upper panels encode the isometry constraints depicted in the lower panels. Specifically, Fig. 2(a) corresponds to Eq. (2) and Fig. 2(b) corresponds to Eq. (3). Note that for convenience we choose a single arrow direction after contracting the legs. In the same spirit we will omit the arrows on physical legs from here on, since we will always choose these to be incoming.

Using these isometric tensors we can construct tensor networks that identically reduce to a single pair of tensors (the “orthogonality center”) upon contraction with its conjugate, which are called isoTNS [13]. In Fig. 3 we show a few snapshots of this reduction for a 2D isoTNS on a 3×3 square

FIG. 2. The isometry constraints are encoded by decorating the tensor legs with arrows (top row). Here panel (a) corresponds to Eq. (2) and panel (b) corresponds to Eq. (3), where in the bottom figures the tensor shown in the top panel is contracted with its conjugate to yield identities.

lattice. Before employing the isometry constraints we have the network depicted in Fig. 3(a). Here we allow the red and gray legs to have distinct bond dimensions χ and D , which is motivated by the observation that for the chosen pattern of arrows we can first reduce the network to a canonicalized MPS amplitude, as shown in Fig. 3(b). Note that the gray legs which stem from the reduction effectively enlarge the local Hilbert space of the MPS.

The distinction between red MPS bonds and gray non-MPS bonds is an important aspect of the isoTNS ansatz, as it will allow us to increase the accuracy of correlators while increasing the bond dimension on only a small part of the tensor network. In Sec. III C we will see that in some situations the accuracy is indeed greatly increased when we increase χ while keeping D constant. This is especially favorable because we will see in Sec. II B 4 that the computational cost of the time-evolution algorithm scales significantly more favorably in χ than in D .

Due to the remaining isometry structure we can further reduce the MPS amplitude down to a single site, as shown in Fig. 3(c). This final site, which is called the orthogonality center and which is colored red in Fig. 3, therefore fully encodes the norm, just like the orthogonality center of MPS in one dimension [4]. It also encodes the one-body correlators of that site, but if we want to calculate its two-body correlator with another site we can already no longer reduce the network to the orthogonality center. In terms of the arrow language this reduction requires that the arrows flow towards the MPS and within the MPS towards the orthogonality center, which can be characterized as having only incoming arrows.

For example, say we want to calculate a local correlator between the two nearest neighbors of the center in Fig. 3(b). In this case the MPS reduction that yields Fig. 3(c) from Fig. 3(b) would be halted at the operators, since there we can no longer utilize the isometry relations from Fig. 2. Consequently we also cannot utilize the isometric tensors that lie between the operators, meaning that we are left with a MPS correlator instead of a single-site correlator, the accuracy of which is controlled by χ . This is illustrated for $\langle \text{isoTNS} | O_1 O_2 | \text{isoTNS} \rangle$ in Fig. 4.

Clearly the calculation of a MPS correlator is more costly than the single-site correlator, but crucially it still scales polynomially in its bond dimension and size [4]. When we instead

consider local correlators with one or more operators located outside of the MPS, the isometry reductions yield a genuine 2D network, which would have to be contracted in order to get the correlator. As a result the correlator can no longer be computed efficiently, illustrating the importance of being able to move the embedded MPS through the network.

It should be noted that by calculating correlators as MPS correlators we do reduce the formal expressibility of the ansatz, as compared to generic PEPS, since it is known that MPS with finite χ can only encode exponentially decaying correlations [4] whereas generic PEPS can also encode algebraic correlations [23]. This means that the isoTNS ansatz can only encode exponential correlations. Nonetheless, for finite systems the bond dimension can always be chosen large enough to encode algebraic correlations which are cut off by the system size.

Before discussing the time-evolution algorithm, which we will do directly for 3D isoTNS, we consider how Fig. 3 generalizes to three dimensions. First we note that the suitable embedding even for two dimensions is actually a tree tensor network (TTN) with the geometry of a star [24,25], which turns into a MPS when the center occupies a corner as in Fig. 3. When it instead occupies the bulk there emerge four MPS strands from the center, which becomes six for a 3D cubic lattice. As long as there are no loops we can put the TTN in canonical form. The principles that underlie the isometry reduction generalize directly to three dimensions, and in Fig. 5 we show the 3D analog of Fig. 3.

B. Evolving 3D isoTNS

We will now explain how the TTN is moved through the 3D isoTNS during its trotterized time evolution, specializing to nearest-neighbor interacting Hamiltonians. In order to take advantage of the isoTNS representation we always apply evolution gates when all sites occupy the TTN, with the orthogonality center at one of the sites. We will see that this gives rise to a threefold-nested TEBD, which we call TEBD³ in analogy to TEBD² for 2D isoTNS.

The evolution takes place columnwise, starting with the middle slice in Fig. 5(a) that is shown separately in Fig. 6(a). To this end we first trotterize the evolution operator $\exp(-d\tau H)$, which is for an imaginary time step of size $d\tau$, in terms of columns c_x , c_y , and c_z . At first order in $d\tau$ we can trotterize as

$$e^{-d\tau H} \approx \prod_{c_x} e^{-d\tau H_{c_x}} \prod_{c_y} e^{-d\tau H_{c_y}} \prod_{c_z} e^{-d\tau H_{c_z}}, \quad (4)$$

where the error is $\mathcal{O}(d\tau^2)$ and stems from the noncommutativity of columns that intersect. With only a bit more effort we can trotterize at second order:

$$e^{-d\tau H} \approx \prod_{c_x} e^{-\frac{d\tau}{2} H_{c_x}} \prod_{c_y} e^{-\frac{d\tau}{2} H_{c_y}} \prod_{c_z} e^{-d\tau H_{c_z}} \prod_{c_y} e^{-\frac{d\tau}{2} H_{c_y}} \prod_{c_x} e^{-\frac{d\tau}{2} H_{c_x}}, \quad (5)$$

which has error $\mathcal{O}(d\tau^3)$. Here subsequent c_x terms can be combined when performing multiple time steps. In

FIG. 3. Multiple stages of the reduction of $(\text{isoTNS} | \text{isoTNS})$ to a contraction of just two tensors, illustrated for a 2D isoTNS on a 3×3 square lattice. Panel (a) shows the amplitude before employing the isometry constraints, where the gray bonds have dimension D , the red bonds have dimension χ , and the blue bonds have dimension d . Panel (b) shows the canonical MPS that is obtained after utilizing a large part of the isometry structure. Panel (c) shows the final pair of contracted tensors that remains after fully utilizing the isometry structure.

Appendix B we investigate the interplay between the trotterization error and the various truncation errors.

Since we will be dealing exclusively with nearest-neighbor interacting systems, each column is further trotterized in terms of two-body gates. At first order this gives us

$$e^{-d\tau H_{\text{cx}}} = \prod_{b_i} e^{-d\tau h_{b_i}} + \mathcal{O}(d\tau^2), \quad (6)$$

where b_i labels the bonds in the column and $h_{b_i} \in \mathbb{C}^{4 \times 4}$ represents the Hamiltonian on b_i . At second order we get

$$e^{-d\tau H_{\text{cx}}} = \prod_{i=1}^L e^{-\frac{d\tau}{2} h_{b_i}} \prod_{i=L}^1 e^{-\frac{d\tau}{2} h_{b_i}} + \mathcal{O}(d\tau^3), \quad (7)$$

which combined with Eq. (5) yields an overall second-order trotterization.

1. Evolving the columns of a slice

To begin the evolution we apply a two-body gate at the initial orthogonality center, as illustrated in Fig. 6(b). For clarity we omit the transverse legs in most of the panels, only restoring them when crucial for interpretation. Before contracting the gate with the tensors in Fig. 7(a) we apply a QR decomposition into an orthogonal matrix Q and a triangular matrix R at each tensor to get a reduced bond [10], yielding the configuration in Fig. 7(b) where we now get to evolve a reduced space because the orange transient bonds are much

smaller than the dark-red bonds (which consist of many virtual legs).

After contracting the gate with the reduced tensors we perform a truncated singular value decomposition (SVD) $A \approx U s V^\dagger$, with U an isometry and $s V^\dagger$ the new orthogonality center, to regain the reduced bond while shifting the orthogonality center, giving the configuration in Fig. 7(c). To ensure that the new bond is not larger than the bond in Fig. 7(a) we often need to truncate the singular values s , which can be done optimally since we are at the orthogonality center. After reabsorbing the reduced tensors we have now shifted the orthogonality center by one site while evolving the bond, giving us the configuration in Fig. 6(c).

We then evolve the next bond, so that we end up at the bottom as shown in Fig. 6(d). To evolve the next column we first need to transfer the TTN strand to this column, for which we use the column-splitting procedure introduced in [13] which is sequence of triangle splittings. This is illustrated in Figs. 6(e)–6(h).

A single triangle splitting is shown in Fig. 8. By performing two truncated SVDs on the orthogonality center in Fig. 8(a) we get the decomposition in Fig. 8(b). To improve the quality of the column splitting, i.e., reduce the information loss in obtaining Fig. 8(g) from Fig. 8(d), we follow [13] and reduce the bipartite entanglement between the right and upper tensors of the triangle in Fig. 8(b). For this we insert a pair of unitary “disentangler” $U^\dagger U = \mathbb{1}$ and optimize them such that the α -Rényi entropy S_α between these tensors is minimized. If we put the orthogonality center s on this bond we can write

$$S_\alpha = \frac{1}{1-\alpha} \ln \sum_i s_i^\alpha, \quad (8)$$

where α is the Rényi order. For $\alpha < 1$ this quantity is known to provide a bound on MPS precision [26]. After minimizing S_α we have a triangle with minimal entanglement across its red bond, and therefore we will end up with a split-off column that has minimal vertical entanglement. The disentangler can be easily optimized with gradient descent. In the case of $\alpha = 2$ there is also a cheaper optimization algorithm [27], but as in [13] we find that $\alpha < 1$ gives the best performance.

FIG. 4. The final product of the isometry reductions when calculating the local two-body correlator $(\text{isoTNS} | O_1 O_2 | \text{isoTNS})$, leaving us with a MPS correlator that has to be fully computed.

FIG. 5. The reduction of a particular 3D isoTNS configuration as it would occur upon contraction with its conjugate. In panel (a) we show the initial network. In panel (b) we show the canonical TTN that would remain after utilizing most of the isometry structure. In panel (c) we show how the canonical TTN further reduces to a single site, i.e., the orthogonality center.

After performing the triangle splitting on the orthogonality center in Fig. 6(d) we absorb the new center upwards as in Figs. 6(e) and 8(c). To continue the column splitting we then temporarily combine the legs as designated in Fig. 8(c) and repeat the triangle splitting, giving the configuration in Fig. 6(f), which becomes Fig. 6(g) after a truncated SVD. This completes the column splitting, and after absorbing the transient column into the next column we have successfully

transferred the TTN strand, as shown in Fig. 6(h). This absorption increases the vertical bond dimension from χ to $D\chi$, which we truncate back to χ before proceeding with the evolution. After repeating the steps in Fig. 6 on the middle column in Fig. 6(h), and subsequently evolving the final column, we have moved the TTN strand from the first to the last column while evolving all columns. The final configuration is shown in Fig. 10(a).

FIG. 6. The evolution of a slice in the 3D isoTNS, where for clarity we omit the transverse legs in most of the panels. In panel (a) we show the initial slice, which is the middle slice from Fig. 5. In panel (b) we apply a two-body gate at the orthogonality center, after which in panel (c) we shift the orthogonality center downwards using a SVD. In panel (d) we repeat this after which the orthogonality center is at the bottom of the first column. In panel (e) we perform the triangle splitting from Fig. 8 on the orthogonality center, which we repeat in panel (f). Now that we have reached the top of the column we perform a truncated SVD to finalize the column splitting, after which we have transferred the TTN strand to a temporary column that has only virtual legs, as shown in panel (g) where we restored the transverse legs to emphasize that the virtual column has no transverse legs. By absorbing the virtual column into the neighboring isometry column, yielding the configuration in panel (h), we have finally moved the TTN strand to the middle column. The increased vertical bond dimension $D\chi$ is truncated back to χ before evolving the new column. Repeating the column splitting and evolution once more we will have evolved all columns in the slice.

FIG. 7. The reduced time evolution used in TEBD³. The bond depicted in panel (a) is to be evolved, to which end we first create a reduced bond by applying a QR decomposition at each tensor. This yields a pair of transient bonds with dimension $d\chi$ that are colored orange. Because these orange bonds are much smaller than the dark-red bonds, the gate application shown in panel (b) and the subsequent truncated SVD that yields panel (c) are much cheaper than when working in the gauge from panel (a). Having evolved the reduced bond we reabsorb the reduced tensors to get the original bond shown in panel (d), for which the orthogonality center has been shifted relative to panel (a).

2. Splitting a slice

Having evolved the columns in the first slice we now want to transfer the TTN strands to the next slice and repeat the evolution. To achieve this we use a tetrahedron splitting on the orthogonality center, which is illustrated in Fig. 9. By using a series of such splittings we can split off a transient slice containing the TTN strands, converting Fig. 10 into Fig. 10(e). Starting at the orthogonality center in Fig. 10(a), shown separately in Fig. 9(a), we perform a sequence of SVDs to get the tetrahedron in Fig. 9(b). The face constituted by the tensors A , B , and C is part of the new slice.

To improve the fidelity of the slice splitting we now need to consider the tripartite entanglement between A , B , and C . Hence we insert a pair of tripartite disentangler $U^\dagger U = \mathbb{1}$

FIG. 8. The triangle splitting of the orthogonality center that is used multiple times in succession to shift a TTN strand to a neighboring column. By performing two truncated SVDs we split the orthogonality center shown in panel (a) into the triangle shown in panel (b), where the upper tensor is the new orthogonality center and where the two lower tensors are isometries. Note that the top and bottom-right tensors have only virtual legs. The amount of truncation during the SVDs is based on the usual color coding, and we have inserted a pair of disentanglers $U^\dagger U = \mathbb{1}$, depicted as yellow cylinders, in order to minimize the entanglement across the red bond. After the splitting we absorb the top tensor upwards, as shown in panel (c). To now perform the splitting from panel (b) on the new orthogonality center we first temporarily combine the tilted legs as indicated by the arrows, so that the orthogonality center again looks as in panel (a).

FIG. 9. The tetrahedron splitting used in the slice splitting that transfers the TTN strands to the next slice. In panel (a) we show the initial orthogonality center. In panel (b) we show the tetrahedron that results from three consecutive SVDs, where the pair of tripartite disentanglers is depicted as a yellow triangle, which reduces the tripartite entanglement $S_{3_\alpha}(A|B|C)$. Note that only one of the four produced tensors has a physical leg. In panel (c) we have absorbed A upwards and B backwards. To repeat the tetrahedron splitting on the new orthogonality center we temporarily combine the tilted legs as indicated by the arrows.

that minimizes a tripartite extension of the α -Rényi entropy S_α (similar to the tripartite information I_3 [28]):

$$S_{3_\alpha}(A|B|C) = S_\alpha(A|BC) + S_\alpha(B|AC) + S_\alpha(C|AB), \quad (9)$$

where, e.g., $S_\alpha(A|BC)$ is the bipartite α -Rényi entropy between tensor A and the contraction of B and C . We minimize $S_{3_\alpha}(A|B|C)$ by iterating over its terms and performing a single step of bipartite disentangling each time. It should be noted that each bipartite disentangler is here a three-body operator. After minimizing $S_{3_\alpha}(A|B|C)$ we get a triangle ABC with minimal entanglement, and hence we will end up with a split-off slice that has minimal internal entanglement.

In Appendix A we compare this tripartite disentangler with a direct 3D extension of the bipartite disentangler from Fig. 8, where we replace each side of the yellow triangle in Fig. 9 by a pair of bipartite disentanglers.

To continue the slice splitting we absorb A upwards and B backwards, which is illustrated in Figs. 9(c) and 10(b). We then temporarily combine the legs as indicated in Fig. 9(c) and repeat the tetrahedron splitting, yielding Fig. 10(c). Note that the new vertical red bond is enlarged to $D\chi$. With the orthogonality center at the top we perform a triangle splitting to get Fig. 10(d), where we also truncated the enlarged vertical bonds back to χ . This produces the first column of the new slice. By combining the tilted legs as in Fig. 9 we can repeat the splitting to move the TTN strand to the final column. We complete the slice splitting by performing the column splitting from Figs. 6(e)–6(h), resulting in Fig. 10(e). The transient slice is then absorbed into the next slice, after which the TTN strands have been successfully transferred, as illustrated in Fig. 10(f). Before continuing with the evolution we truncate the enlarged bonds.

3. Evolving the full network

With the machinery developed in the previous sections we can now perform a trotterized time evolution where all truncation occurs at the orthogonality center, so that the local evolution is globally optimal. To illustrate this we consider the details of a single TEBD³ iteration, starting from the network in Fig. 11(a).

FIG. 10. The slice splitting by which we transfer the two TTN strands to the next slice, after evolving the vertical bonds of a slice. In panel (a) we show the initial isoTNS, which is obtained from Fig. 5 by evolving the middle slice as in Fig. 6. In panel (b) we show the first step of the slice splitting, where we employ the tetrahedron splitting from Fig. 9 on the orthogonality center, thereby moving it upwards by one site. After temporarily combining the tilted legs with the horizontal legs as indicated by the arrows in Fig. 9(c) we repeat the tetrahedron splitting to get panel (c). Finally we apply the triangle splitting from Fig. 8 to the uppermost tensor in panel (d). We then move the orthogonality center to the bottom of the column, resulting in panel (d). We repeat this combination of evolving and column splitting until we reach the final column, which is instead split via the procedure described in Figs. 6(e)–6(h), yielding the split-off slice in panel (e). To finalize the transfer we absorb the transient slice into the neighboring slice, yielding panel (f). The increased bond dimensions are truncated before continuing the evolution.

FIG. 11. A single iteration of TEBD^3 . Starting from the initial configuration in panel (a) we evolve all columns to get panel (b). Here the evolved gray bonds are colored green and the evolved red bonds are colored dark green. To evolve the next set of columns we rotate the network as indicated by the numbering of the corners, yielding panel (c), which becomes panel (d) after evolving the new columns. To evolve the final set of columns we rotate the network to panel (e), which turns into panel (f) upon evolving the columns. With all bonds evolved this concludes a single iteration of TEBD^3 .

We evolve the columns of the first slice, following the procedure in Fig. 6, and subsequently move the TTN strands to the neighboring slice using the slice splitting of Fig. 10. Repeating these operations on the middle slice and afterwards evolving the final slice we end up with Fig. 11(b), where all columns have now been evolved as indicated by the green bonds.

To evolve the rest of the bonds we rotate the network in Fig. 11(b) such that the TTN again has the position from Fig. 11(a) while the columns consist of nonevolved bonds. We have numbered the corners in Fig. 11 to show a possible way of doing this. After evolving the columns, yielding Fig. 11(d), we rotate the network to get Fig. 11(e), which becomes Fig. 11(f) upon evolving the final columns. Since all bonds have now been evolved we have completed the iteration of TEBD^3 .

The TEBD^3 algorithm has two main sources of error: the error due to the multitude of truncated SVDs and the error due to the trotterization. In [13] it was found that for TEBD^2 the column-splitting truncation error and trotterization error conspire to yield an energy minimum in $d\tau$ space. In Appendix B we show that the same occurs for TEBD^3 , now resulting from the interplay between the column- and slice-splitting truncation errors and the trotterization error.

4. Computational complexity

The computational complexity of the TEBD^3 algorithm can be attributed to various operations, in particular to the SVDs that occur during these operations. Here we include only terms that are potentially leading in either d , D , or χ . We moreover assumed that we perform full SVDs instead of partial SVDs in obtaining all estimates.

Starting with the slice splitting and subsequent absorption, we find that the tetrahedron splitting has cost $\mathcal{O}[dD^{10}\chi^2 \min(d, D\chi^2)]$ when it is performed on a tensor with the maximum amount of combined legs (see Fig. 9). The truncation of the enlarged bonds after absorbing the split-off slice has cost $\mathcal{O}(dD\chi^6)$.

Other contributions arise from the evolution of the bonds in a column, where the pairs of QR decompositions that precede gate application have cost $\mathcal{O}[d\chi^5 \min(d, \chi^2)]$. These decompositions yield pairs of intermediate bonds with sizes η and ζ , which depend on the external bonds involved in the decompositions, so that the subsequent SVD has cost $\mathcal{O}[d^3 \min(\eta\zeta^2, \eta^2\zeta)]$. The pair that potentially yields the leading order in D occurs in the bulk and has $\eta = \zeta = \chi \min(d, D^4)$.

When simultaneously $d \leq \chi^2$ and D^4 the total cost reduces to $\mathcal{O}(d^2 D^{10} \chi^2) + \mathcal{O}(dD\chi^6) + \mathcal{O}(d^6 \chi^3)$. Because the costliest operations are performed roughly N times we have the linear scaling $\mathcal{O}(N)$ in system size.

III. BENCHMARKING

A. The benchmarking system

As a proof of principle we probe the accuracy of TEBD^3 for imaginary time evolution to find an isoTNS approximation for the many-body ground state of a simple 3D quantum many-body system: the ferromagnetic transverse field Ising

model (TFIM) on a cubic lattice with OBC:

$$H = - \sum_{\langle ij \rangle} \sigma_i^z \sigma_j^z - h \sum_i \sigma_i^x. \quad (10)$$

Here $\vec{\sigma}_i \equiv (\sigma_i^x, \sigma_i^y, \sigma_i^z)$ corresponds to a spin 1/2 on site i , with $\sigma^{x,y,z}$ the usual Pauli matrices, and the two-body sum runs over nearest-neighbor pairs $\langle i, j \rangle$.

Working in the σ^z basis, we see from Eq. (10) that the TFIM Hamiltonian becomes classical for $h \rightarrow 0$ where it has a twofold degenerate ferromagnetic ground state $|\uparrow\uparrow\ldots\uparrow\rangle$ and $|\downarrow\downarrow\ldots\downarrow\rangle$. Quantum fluctuations are generated by the uniform magnetic transverse (i.e., along the x direction) field of strength h . In the limit of strong fields $h \rightarrow \infty$, the ground state is unique and aligned with the field and a simple product state $|\rightarrow\rightarrow\rightarrow\ldots\rightarrow\rangle$ in the σ_x basis. Between these two limits, the competition between the x and z basis leads to a much more complex ground state [29] and to a quantum phase transition as a function of h between the ferromagnetic phase at $h < h_c$ and the polarized phase at $h > h_c$. The value of the critical field on the 3D simple cubic lattice was numerically estimated to be $h_c \approx 5.158\,13(6)$ [30].

Our main reason for choosing this model as a benchmark is its simplicity and the fact that it can be solved exactly using quantum Monte Carlo using the stochastic series expansion [31,32]. Our two benchmark observables are the energy density $E/N = \langle H \rangle / N$ and the x magnetization $m_x = \sum_i \langle \sigma_i^x \rangle / N$, which we compare to exact values in order to assess the accuracy of TEBD^3 .

Quantum phase transitions are associated with a divergent correlation length and are therefore extremely challenging to study using tensor network methods. While a phase transition occurs at a singular point in the thermodynamic limit, a critical region is expected for finite systems. In the case of OBC the region gets shifted towards $h < h_c$, which is especially pronounced for the smaller L , where the ratio of one- to two-body couplings is significantly larger. We expect the 3D isoTNS ansatz to be particularly challenged in this critical region, whereas its performance is likely to improve when progressing deeper into both phases, since the ground state in both phases is a dressed product state.

B. The benchmarking procedure

We perform imaginary time TEBD^3 propagation of the wave function represented by an isoTNS with bond dimensions (D, χ) for the TFIM (10) with $L = 3, 4, 10$ at various fields $h = 0.5, 1.0, \dots, 8$ (across the critical point). For $L = 3$ we use bond dimensions $D = 2, 4, 6$ with $\chi = 2D, 4D, 6D$, which allows us to observe the convergence of both E/N and m_x towards the exact values. For $L > 3$ we are more constrained in our choice of (D, χ) due to the large cost as derived in Sec. II B 4. Hence for $L = 4$ we use $D = 2, 3$ with $\chi = 2D, 4D, 6D$, and to illustrate the capacity of TEBD^3 to reach large system sizes we consider a $(2,4)$ isoTNS for $L = 10$.

For the initial state we take a σ_z product state $|\psi_{\text{ini}}\rangle = |\uparrow\uparrow\ldots\uparrow\rangle$, which we evolve in imaginary time until the energy density is converged. This means that we perform the opera-

tion

$$|\psi_0\rangle \approx e^{-\beta \hat{H}} |\uparrow\uparrow \dots \uparrow\rangle, \quad (11)$$

which for TEBD algorithms is done by trotterizing it into small time steps $d\tau$ as explained in Sec. II B. We typically evolve to $\beta \gtrsim 40$.

As mentioned Sec. II B 3, TEBD³ has an energy minimum in $d\tau$ space, so for each simulation we choose $d\tau$ such that it coincides with the minimum. In Fig. 17 of Appendix B we show the $d\tau$ minima for various (D, χ) on a $L = 3$ lattice in the critical region. Similarly, we consider various Rényi entropy orders α for the disentangling and pick the one with lowest energy. It is usually between $\alpha = 1/2$ and 1, although it varies across the phase diagram.

The exact values for the energy density were obtained with the stochastic series expansion, for which we used the ALPS library [33,34]. Here convergence to the ground state was checked by a β -doubling scheme. For $L = 3$ we also performed Lanczos exact diagonalization, providing exact energy density and additionally the x magnetization of the ground state at all h .

In order to provide a reference for the quality of our TEBD³ results, we furthermore performed TEBD² calculations for the 2D TFIM using the algorithm detailed in [13]. Here we chose a comparable system size of 5×5 , for which exact Lanczos values are easily obtained.

C. Results

In this section we present the results of our TEBD³ benchmarks for the 3D TFIM ground state.

We start off with $L = 3$, for which the performance can be comprehensively probed. In Fig. 12 we have plotted E/N and m_x across a range of h . Here the thermodynamic critical point is denoted by a gray dotted line, but as mentioned in Sec. III A this point is spread into an extended region for finite systems, and furthermore shifted to smaller $h < h_c$ due to the use of OBC.

The top panel of Fig. 12 shows E/N across a range of h for various (D, χ) , along with the exact Lanczos values as a red dotted line. The inset shows the relative error in E/N . Here we see that deep in the polarized and ferromagnetic phases the accuracy is high, whereas in the critical region around $h \approx 3$ the performance is clearly worse.

For $D = 2$ and 4 we see that the accuracy is converged in χ , meaning that D must be increased (i.e., the column and slice splitting need higher fidelity) to further improve the performance. For $D = 6$ it is not clear whether convergence is reached, meaning that convergence becomes slower as D grows. On a related note, it can be seen that (4,16) performs better than (6,12), illustrating why large D is only useful when combined with a few multiples larger χ .

The bottom panel of Fig. 12 shows m_x for various h , again together with the exact Lanczos values as a red dashed line. The inset contains the absolute error $\delta = m_x^e - m_x$ relative to the exact value m_x^e , which is chosen over the relative error since m_x^e vanishes as $h \rightarrow 0$. Deep in the polarized and ferromagnetic phases we see excellent agreement already for small bond dimensions, with the critical region clearly requiring larger tensors to get near the exact line.

FIG. 12. Top: The energy density E/N and its accuracy relative to the exact Lanczos value for the $3 \times 3 \times 3$ cubic TFIM with OBC at various field strengths h across the critical point. Each curve belongs to a different set of bond dimensions (D, χ) , and in the main plot we show the exact values with a red dotted line. Bottom: The x magnetization m_x and its error $\delta = m_x^e - m_x$ relative to the exact Lanczos value m_x^e for various field strengths around the critical point, for the same TFIM system as in the top panel.

Next we consider $L = 4$, for which we show the results in Fig. 13. The curves display similar behavior as for $L = 3$, but due to the high computational cost we did not probe its performance beyond $D = 3$. A noticeable difference is the sharper peak in energy accuracy, which is likely due to the shrinking critical region (that has moreover shifted to larger h). For $L = 4$ the exact values were obtained via Monte Carlo, with the accompanying errors falling inside the linewidth.

Now we consider $L = 10$ (i.e., $N = 10^3$) in Fig. 14. Here we were not able to go beyond $D = 2$, which embodies the large difference in computational complexity between the bond dimensions and system size, as found in Sec. II B 4. We again observe a peak in relative error of E/N at the critical region, which is now centered on the thermodynamic h_c . The x magnetization is now also seen to saturate around h_c .

Overall, the comparison of our results across different system sizes reveals an excellent representation of the ground-state wave function deep in the ferromagnetic phase, even for the relatively small (D, χ) which are reachable at large system sizes on current computers. The accuracy is slightly worse in the polarized phase, especially for the smaller (D, χ) , but nonetheless high accuracies can be reached on relatively

FIG. 13. Top: The energy density E/N and its accuracy relative to the exact QMC value for the $4 \times 4 \times 4$ cubic TFIM with OBC at various field strengths h across the critical point. Each curve belongs to a different set of bond dimensions (D, χ) , and in the main plot we show the exact values with a red dotted line. Bottom: The x magnetization m_x for various field strengths around the critical point, for the same TFIM system as in the top panel.

FIG. 14. Top: The energy density E/N of the (2,4) isoTNS and its accuracy relative to the exact QMC value for the $10 \times 10 \times 10$ cubic TFIM with OBC at various field strengths h across the critical point. The exact values are shown as a red dotted line. Bottom: The x magnetization m_x for various field strengths around the critical point, for the same TFIM system as in the top panel.

IV. CONCLUSION

small lattices. As expected, the accuracy is worst in the critical region, but also here the performance is significantly improved upon increasing the bond dimensions, showing a clear trend toward the exact values both for the energy density and x magnetization.

In order to put the TEBD^3 benchmarks into perspective we have also performed a TEBD^2 benchmark for the 2D TFIM on a 5×5 square lattice. It is clear that the simulation of a 2D system with TEBD^2 is easier than that of a 3D system with TEBD^3 , since the 3D version involves more truncation. We therefore expect TEBD^2 to have higher accuracy for similar bond dimensions.

In Fig. 15 we show E/N and m_x for multiple (D, χ) at various h for the 5×5 TFIM. In the top panel we see that the relative error of E/N is again smallest when deep in the ferromagnetic and polarized phases, with the critical region around $h_c \approx 3$ again posing the biggest challenge. As expected, we see that (4,16) performs better in two dimensions than in three dimensions, reaching just below 0.5% compared to just below 1% in Fig. 12. The same is apparent from the bottom panel, where we see that (4,16) already closely matches the exact m_x in the critical region, which is reached only with (6,36) in Fig. 12.

We have introduced a method for the simulation of 3D quantum lattice models using a representation of the wave function as a 3D isometric tensor network state (isoTNS). Generalizing the method for two dimensions presented in [13] we introduced a tetrahedral splitting and accompanying tripartite disentangling, such that optimal time-evolving block decimation can be carried out in cubic 3D networks. We call the resulting evolution algorithm TEBD^3 .

Our systematic benchmark for the 3D transverse field Ising model in the full range of transverse fields across the critical point reveals that our method yields accurate results, and that the systematic error incurring from finite bond dimensions can be controlled systematically by increasing (D, χ) . This behavior is identical to what is observed in TEBD^2 . The regime close to the critical point is particularly challenging and requires larger bond dimensions, beyond the capacity of our computers for large systems.

While imaginary time evolution using TEBD^3 is arguably the simplest method for finding the ground state of a quantum many-body system, it is known even in one dimension that it is not optimal and that local variational energy minimization (e.g., DMRG in one dimension) is far more efficient. We expect a similar behavior for isoTNS in higher dimensions and it is possible that our results can be further improved by

FIG. 15. Top: The energy density E/N and its accuracy relative to the exact Lanczos value for the 5×5 square TFIM with OBC at various field strengths h across the critical point. Each curve belongs to a different set of bond dimensions (D, χ) , and in the main plot we show the exact values with a red dotted line. Bottom: The x magnetization m_x and its error $\delta = m_x^e - m_x$ relative to the exact Lanczos value m_x^e for various field strengths around the critical point, for the same TFIM system as in the top panel.

the formulation of a DMRG analog for 3D isoTNS, a direction which we leave for future study.

ACKNOWLEDGMENTS

We are grateful to M. Zaletel and F. Pollmann for useful suggestions to further improve isoTNS. We acknowledge financial support from the Deutsche Forschungsgemeinschaft through SFB1143 (Project number 247310070).

APPENDIX A: BIPARTITE VERSUS TRIPARTITE DISENTANGLING

To quantify the difference in performance between the bipartite and tripartite 3D disentanglers from Sec. II B 2 we use each to calculate the ground-state energy density of the $4 \times 4 \times 4$ TFIM at various h . In Fig. 16 we show the performance of the isoTNS configurations (2,12) and (3,15). Here it is clear that the improvement in accuracy when using tripartite over bipartite disentanglers is largest in the critical region $h \in [3, 5]$, whereas the improvement becomes progressively smaller outside of this region.

FIG. 16. A comparison of the bipartite and tripartite disentanglers, showing the energy accuracy for the $4 \times 4 \times 4$ TFIM ground state at various h and two (D, χ) . We see that the improvement in using a tripartite over a bipartite disentangler is most significant in the critical region.

Because the chosen (D, χ) are converged in χ we cannot improve the bipartite curves by further increasing the TTN's bond dimension. This illustrates that the quality of the tetrahedron splitting, and hence of the slice splitting, cannot be compensated by only improving the TTN. This is easily understood if we recognize that the slice splitting is a major component of TEBD³, since it serves to transfer the TTN to the next slice with minimal information loss, and that this part of the algorithm is mainly controlled by D and not χ (see Fig. 9). In particular, even though a larger χ might improve the time evolution on a fixed TTN, this gain is lost when transferring the TTN.

APPENDIX B: MINIMA IN $d\tau$ SPACE

As noted in [13], the error due to the multitude of truncated SVDs and the error due to the trotterization combine to yield an energy minimum in $d\tau$ space for TEBD². Here we show that this also occurs for TEBD³ and we will furthermore illustrate the difference between first- and second-order trotterization.

In Fig. 17 we show the relative error in energy density for the ground state of the $3 \times 3 \times 3$ TFIM at $h = 3.5$, at multiple points in $d\tau$ space. For each (D, χ) we plot both

FIG. 17. The $d\tau$ dependency of the relative error in energy density of the $3 \times 3 \times 3$ TFIM at $h = 3.5$, for various isoTNS configurations (D, χ) .

the first-order ($n = 1$) and second-order ($n = 2$) trotterization results, from which we see that for all considered cases the $n = 2$ minimum lies below the $n = 1$ minimum. We can also

see that the minima shift to lower $d\tau$ as we increase D , for both $n = 1$ and 2, which is in accordance with the findings in [13] for TEBD².

-
- [1] M. B. Hastings, An area law for one-dimensional quantum systems, *J. Stat. Mech.: Theory Exp.* (2007) P08024.
- [2] M. B. Hastings, Locality in Quantum and Markov Dynamics on Lattices and Networks, *Phys. Rev. Lett.* **93**, 140402 (2004).
- [3] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, *Rev. Mod. Phys.* **82**, 277 (2010).
- [4] U. Schollwoeck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (NY)* **326**, 96 (2011).
- [5] S. R. White, Density Matrix Formulation for Quantum Renormalization Groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [6] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions, [arXiv:cond-mat/0407066](https://arxiv.org/abs/cond-mat/0407066).
- [7] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Computational Complexity of Projected Entangled Pair States, *Phys. Rev. Lett.* **98**, 140506 (2007).
- [8] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Unifying projected entangled pair state contractions, *New J. Phys.* **16**, 033014 (2014).
- [9] F. Verstraete, V. Murg, and J. I. Cirac, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, *Adv. Phys.* **57**, 143 (2008).
- [10] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Algorithms for finite projected entangled pair states, *Phys. Rev. B* **90**, 064425 (2014).
- [11] G. Vidal, Efficient Classical Simulation of Slightly Entangled Quantum Computations, *Phys. Rev. Lett.* **91**, 147902 (2003).
- [12] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwoeck, and C. Hubig, Time-evolution methods for matrix-product states, *Ann. Phys. (NY)* **411**, 167998 (2019).
- [13] M. P. Zaletel and F. Pollmann, Isometric Tensor Network States in Two Dimensions, *Phys. Rev. Lett.* **124**, 037201 (2020).
- [14] M. Zaletel, Isometric Tensor Network States in Two Dimensions, lecture delivered at Perimeter Institute, Waterloo, ON, 2019, <https://www2.perimeterinstitute.ca/videos/isometric-tensor-network-states-two-dimensions>.
- [15] R. Haghshenas, M. J. O'Rourke, and G. K.-L. Chan, Conversion of projected entangled pair states into a canonical form, *Phys. Rev. B* **100**, 054404 (2019).
- [16] K. Hyatt and E. M. Stoudenmire, DMRG Approach to Optimizing Two-Dimensional Tensor Networks, [arXiv:1908.08833](https://arxiv.org/abs/1908.08833).
- [17] T. Soejima, K. Siva, N. Bultinck, S. Chatterjee, F. Pollmann, and M. P. Zaletel, Isometric tensor network representation of string-net liquids, *Phys. Rev. B* **101**, 085117 (2020).
- [18] P. Teng, Generalization of the tensor renormalization group approach to 3-d or higher dimensions, *Physica A* **472**, 117 (2017).
- [19] A. García-Sáez and J. I. Latorre, Renormalization group contraction of tensor networks in three dimensions, *Phys. Rev. B* **87**, 085130 (2013).
- [20] S. Anders, M. B. Plenio, W. Dür, F. Verstraete, and H.-J. Briegel, Ground-State Approximation for Strongly Interacting Spin Systems in Arbitrary Spatial Dimension, *Phys. Rev. Lett.* **97**, 107206 (2006).
- [21] A. Sfondrini, J. Cerrillo, N. Schuch, and J. I. Cirac, Simulating two- and three-dimensional frustrated quantum systems with string-bond states, *Phys. Rev. B* **81**, 214426 (2010).
- [22] M. Schmitt and M. Heyl, Quantum dynamics in transverse-field Ising models from classical networks, *SciPost Phys.* **4**, 013 (2018).
- [23] 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).
- [24] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, *Phys. Rev. A* **74**, 022320 (2006).
- [25] F. A. Y. N. Schroeder, D. H. P. Turban, A. J. Musser, N. D. M. Hine, and A. W. Chin, Tensor network simulation of multi-environmental open quantum dynamics via machine learning and entanglement renormalisation, *Nat. Commun.* **10**, 1062 (2019).
- [26] F. Verstraete and J. I. Cirac, Matrix product states represent ground states faithfully, *Phys. Rev. B* **73**, 094423 (2006).
- [27] J. Hauschild, E. Leviatan, J. H. Bardarson, E. Altman, M. P. Zaletel, and F. Pollmann, Finding purifications with minimal entanglement, *Phys. Rev. B* **98**, 235163 (2018).
- [28] M. Rangamani and M. Rota, Entanglement structures in qubit systems, *J. Phys. A: Math. Theor.* **48**, 385301 (2015).
- [29] D. J. Luitz, F. Alet, and N. Laflorencie, Universal Behavior Beyond Multifractality in Quantum Many-Body Systems, *Phys. Rev. Lett.* **112**, 057203 (2014).
- [30] H. W. J. Blöte and Y. Deng, Cluster Monte Carlo simulation of the transverse Ising model, *Phys. Rev. E* **66**, 066110 (2002).
- [31] A. W. Sandvik and J. Kurkijärvi, Quantum Monte Carlo simulation method for spin systems, *Phys. Rev. B* **43**, 5950 (1991).
- [32] A. W. Sandvik, A generalization of Handscomb's quantum monte carlo scheme-application to the 1D Hubbard model, *J. Phys. A: Math. Gen.* **25**, 3667 (1992).
- [33] B. Bauer, L. D. Carr, H. G. Evertz, A. Feiguin, J. Freire, S. Fuchs, L. Gamper, J. Gukelberger, E. Gull, S. Guertler, A. Hehn, R. Igarashi, S. V. Isakov, D. Koop, P. N. Ma, P. Mates, H. Matsuo, O. Parcollet, G. Pawłowski, J. D. Picon *et al.*, The ALPS project release 2.0: Open source software for strongly correlated systems, *J. Stat. Mech.: Theory Exp.* (2011) P05001.
- [34] A. F. Albuquerque, F. Alet, P. Corboz, P. Dayal, A. Feiguin, S. Fuchs, L. Gamper, E. Gull, S. Gürtler, A. Honecker, R. Igarashi, M. Körner, A. Kozhevnikov, A. Läuchli, S. R. Manmana, M. Matsumoto, I. P. McCulloch, F. Michel, R. M. Noack, G. Pawłowski *et al.*, The ALPS project release 1.3: Open-source software for strongly correlated systems, *J. Magn. Magn. Mater.* **310**, 1187 (2007).