

Simulating both parity sectors of the Hubbard model with tensor networks

Manuel Schneider^{1,2,*}, Johann Ostmeyer^{3,4}, Karl Jansen¹, Thomas Luu^{3,5,6} and Carsten Urbach^{3,4}

¹*NIC, DESY Zeuthen, Platanenallee 6, 15738 Zeuthen, Germany*

²*Institut für Physik, Humboldt-Universität zu Berlin, Newtonstraße 15, 12489 Berlin, Germany*

³*Helmholtz-Institut für Strahlen- und Kernphysik, University of Bonn, Nussallee 14-16, 53115 Bonn, Germany*

⁴*Bethe Center for Theoretical Physics, University of Bonn, Nussallee 12, 53115 Bonn, Germany*

⁵*Institute for Advanced Simulation, Forschungszentrum Jülich, 54245 Jülich, Germany*

⁶*Institut für Kernphysik, Forschungszentrum Jülich, 54245 Jülich, Germany*



(Received 9 July 2021; accepted 2 September 2021; published 11 October 2021)

Tensor networks are a powerful tool to simulate a variety of different physical models, including those that suffer from the sign problem in Monte Carlo simulations. The Hubbard model on the honeycomb lattice with nonzero chemical potential is one such problem. Our method is based on projected entangled pair states using imaginary-time evolution. We demonstrate that it provides accurate estimators for the ground state of the model, including cases where Monte Carlo simulations fail miserably. In particular, it shows near to optimal, that is linear, scaling in lattice size. We also present an approach to directly simulate the subspace with an odd number of fermions. It allows to independently determine the ground state in both sectors. Without a chemical potential, this corresponds to half-filling and the lowest-energy state with one additional electron or hole. We identify several stability issues, such as degenerate ground states and large single-particle gaps, and provide possible fixes.

DOI: [10.1103/PhysRevB.104.155118](https://doi.org/10.1103/PhysRevB.104.155118)

I. INTRODUCTION

Projected entangled pair states (PEPS) are a higher-dimensional generalization of matrix product states (MPS) [1] and have successfully been used to approximate various physical systems. Examples include various two-dimensional (2D) classical and quantum spin models [2,3], spin liquids [4,5], as well as select lattice non-Abelian gauge theories [6,7]. Although originally developed for bosonic systems, calculations nowadays readily use fermionic PEPS [8] to investigate systems with fermionic degrees of freedom.

As the method is variational and deterministic, calculations using tensor networks such as PEPS do not suffer from numerical issues seen in stochastic simulations when dealing with systems with nonzero chemical potential μ . In such cases, an induced numerical sign problem can completely preclude a stochastic simulation, while tensor networks instead are completely immune and provide results equally precise as compared to cases when $\mu = 0$. Furthermore, tensor network calculations of global ground-state energies of various systems are achieving remarkable accuracy [9]. With continual advancements in algorithmic efficiency, tensor networks are fast becoming the “go-to” method for extracting the global ground-state properties of low-dimensional systems.

The situation is less clear, however, when dealing with excited states, or ground-state energies in sectors of opposite parity to that of the global ground state. To date, tensor networks could only access a few low-lying energy states by calculating the ground state first and projecting to an

orthogonal subspace in the following iterations [10]. This is unfortunate, as an understanding of the spectrum of states *relative* to the global ground state provides information on novel forms of collective behavior of the system, such as spin and antiferromagnetic correlations [11], or induced Mott transitions [12], to name a few examples. A quintessential example in this regard is the Hubbard model in two dimensions [13,14]. On an infinite honeycomb, or hexagonal lattice, for example, the Hubbard model exhibits a quantum phase transition from a semimetal to insulator state at a particular critical coupling [15,16]. Such a transition can be ascertained by comparing the ground-state energy of the odd-parity sector to that of the even-parity sector, as was done in [17]. Numerical investigations of the Hubbard model have mostly been confined to chemical potential $\mu = 0$ so as to avoid the sign problem, which means that much of its phase diagram is not known. Thus, were tensor networks able to obtain information of odd-parity states with equal precision as obtained with the global ground state and with arbitrary μ , a new era in simulations of nonperturbative phenomena would occur for low-dimensional systems.

In this paper we build off the formalism presented in [8] and apply it to the odd-parity sector of the hexagonal Hubbard model as a test case. Although tensor networks have been applied to the Hubbard model using infinite PEPS (iPEPS) [18], such investigations analyzed the global ground state of the Hubbard model on a the square lattice. In our case, in addition to considering the hexagonal lattice which has important applications in the study of graphene and its derivatives [19,20], we extract the lowest-energy state in the odd-parity sector of this system, giving us information about the single-particle gap. We perform these calculations for select values of onsite

*manuel.schneider@desy.de

Hubbard term U and chemical potential μ , benchmarking our results to exact results where possible. For systems too large for comparison with exact calculations, our results represent predictions.

Our method for simulating in opposite parity sectors is not a complete panacea for tensor network calculations of excited states. Indeed, we encounter various numerical difficulties and limitations which we carefully document and, where possible, address. Still, our findings provide proof of principle that excited states can be readily extracted using tensor networks, and we anticipate continued development along the lines of our research.

Our paper is organized as follows: In the following section we give a cursory overview of the PEPS formalism and its connection to imaginary-time evolution. We present our algorithm for fermions that allows us to simulate in opposite parity sectors. We discuss how we implement various update schemes and demonstrate with hard examples the expected scaling of our algorithm with system size, bond dimension, and boundary MPS dimension. We also introduce the Hubbard Hamiltonian in this section and explain how it is adapted to our method. In Sec. III we enumerate our various improvements to our algorithm to accelerate convergence as well as increase precision. We take pains to enumerate the limitations of our procedure so as to present to the reader a clear and honest assessment of our method's efficacy. Where possible, we provide explanations for the sources of these limitations and potential resolutions. We then show our results for the Hubbard model in Sec. IV. Our results demonstrably show the ability of our method to simulate both parity sectors of the problem, obtaining precise results for both ground states in both cases and therefore gaining information about the single-particle gap. Finally, we recapitulate in Sec. V.

II. FORMALISM

A. PEPS

A projected entangled pair state (PEPS) [21,22] is an ansatz for a physical state. It is a versatile tool to incorporate different entanglement structures of two- and higher-dimensional systems and especially applicable for ground-state simulations.

We want to study a system of N lattice sites on a $(L_x \times L_y)$ grid,¹ where the local dimension of the Hilbert space is d . Representing a state would therefore require d^N complex numbers, which is not feasible to store on a computer system for larger lattices. Therefore, a truncation of the Hilbert space is needed. A PEPS reduces the number of parameters drastically while still giving a good approximation of specific states. The truncation corresponds to a cutoff in the entanglement entropy. It was shown that ground states of local Hamiltonians with an energy gap fulfill an area law and can therefore be well approximated by tensor networks [23]. This holds in one spatial dimension, but tensor networks and PEPS in particular were also applied successfully in higher dimensions [24].

To construct a PEPS, one places a tensor on each lattice site, pictured as circles in Fig. 1. Each tensor carries a physical




FIG. 1. Ket state of a PEPS for a 3×4 fermionic honeycomb lattice. Description of symbology (see text for more details): circles, PEPS tensors; dashed lines, physical indices; solid lines, internal indices; dotted line, parity index; diamonds, swap gates.

index running from 1 to d (dashed lines). Additional indices connect the tensor to its nearest neighbors with indices (links) running from 1 to D_k , where k enumerates the link between the two nearest neighbors (solid lines). Lines connecting two tensors require summation (contraction) over the corresponding index, while open legs denote the open indices of the resulting tensor after contraction of the network. Figure 2 shows different projections of the bra state corresponding to the ket state in Fig. 1. The swap gates and the parity index in Figs. 1 and 2 will be explained in Sec. II C.

In the limit of large D_k this ansatz becomes exact. For larger systems, a truncation bond dimension D is chosen and the indices are bound by $D_k \leq D \forall k$. The accuracy of the ansatz can be improved by increasing D .




FIG. 2. Bra state of a fermionic PEPS. Symbology: circles without stars, conjugated PEPS tensors; dashed lines, physical indices; solid lines, internal indices; dotted line, parity index; diamonds, swap gates. (a) Bra tensors for a 3×4 honeycomb lattice. (b) Definition of starred tensors with four internal indices. Some indices and the corresponding SWAP gates can be removed to obtain the honeycomb tensors.

¹This corresponds to $(L_x/2 \times L_y)$ unit cells of the honeycomb lattice, up to boundary effects.

B. Imaginary-time evolution

Expressing the state in the form of a PEPS allows us to evolve the system in imaginary time to the ground state $|\Psi_0\rangle$ of a given Hamiltonian H . We start with a (random) initial state $|\Psi(0)\rangle$ and evolve it to find the ground state:

$$|\Psi_0\rangle = \lim_{t \rightarrow \infty} \frac{|\Psi(t)\rangle}{\sqrt{\langle \Psi(t) | \Psi(t) \rangle}}, \quad |\Psi(t)\rangle = e^{-Ht} |\Psi(0)\rangle. \quad (1)$$

The time-evolution operator

$$U(t) := e^{-Ht} = (e^{-H\delta t})^m \quad (2)$$

is decomposed into m time slices of length $\delta t = t/m$. Every single time slice is again split into a product of local terms. We focus on a Hamiltonian that consists only of nearest-neighbor interaction terms H_i :

$$H = \sum_{i=1}^{\#H_i} H_i, \quad (3)$$

$$U(\delta t) = \prod_i e^{-H_i \delta t} + O(\delta t^2), \quad (4)$$

$$U(\delta t) = \prod_{i=1 \dots \#H_i} e^{-H_i \frac{\delta t}{2}} \prod_{j=\#H_i \dots 1} e^{-H_j \frac{\delta t}{2}} + O(\delta t^3). \quad (5)$$

Here we used the Suzuki-Trotter expansion [25,26] of first order in Eq. (4) and of second order in Eq. (5). For our calculations we applied the second-order expression, where the ordering of terms in the second product is inverse to the ordering in the first.

This way the time-evolution operator $U(t)$, which acts *globally*, can be decomposed into a sequence of local terms $e^{-H_i \delta t}$. These can be efficiently applied to the PEPS, which we discuss in Sec. IID.

C. Modifications for fermions

If we want to consider fermions in our model, we have to take into account their anticommuting nature. We need to order the fermionic degrees of freedom and make sure that the sign flips once we exchange two fermionic creation or annihilation operators. In Fig. 1 the two-dimensional network is arranged in a way that we can number the physical indices from left to right. The state would be created by applying creation operators in the order of this numbering to the vacuum state. We can apply a nearest-neighbor operator in the y direction directly because the corresponding links have consecutive numbers. Since any observable has an even number of creation and annihilation operators, all sign flips come in pairs and cancel each other. If the creation and annihilation operators of an observable sit on sites with nonconsecutive numbers, however, we would have to take into account all indices with numbers in-between. This happens for the application of a gate in the x direction in our numbering scheme. This would turn nearest-neighbor operators into nonlocal ones. Fortunately, we can avoid this issue by using the scheme developed in [8,27], which we now discuss.

1. Fermion SWAP gate

We implement the fermionic anticommutation relations at the PEPS level rather than at the operator level only. In order

to do this, we first introduce the concept of fermion-number parity. An even number of fermions comes with parity $p = 1$, an odd number with parity $p = -1$, i.e., $p(|0\rangle) = p(|\uparrow\downarrow\rangle) = 1$ and $p(|\uparrow\rangle) = p(|\downarrow\rangle) = -1$. In addition, we assign the parity

$$p = (\underbrace{1, \dots, 1}_{D_e}, \underbrace{-1, \dots, -1}_{D_o}) \quad (6)$$

with $D_e + D_o = D_k$ to every internal tensor index k . Parity conservation requires fermions to be created and annihilated in pairs. We can ensure this on the level of the tensor network by the rule

$$T_{i_1, i_2, \dots, i_r} = 0 \quad \text{if} \quad p(i_1)p(i_2) \dots p(i_r) = -1. \quad (7)$$

The crucial difference to bosonic tensor networks, however, is the aforementioned anticommutation of the creation operators. An exchange of two odd numbers of fermionic operators yields a minus sign. In the PEPS formalism this means that every crossing of two lines introduces additional signs. A minus sign occurs at every crossing for the indices with twice negative parity. In [27] this is achieved by introducing fermionic SWAP gates

$$X_{j_1 j_2}^{i_1 i_2} = \delta_{i_1 j_1} \delta_{i_2 j_2} S(i_1, i_2), \quad (8)$$

$$S(i_1, i_2) = \begin{cases} -1 & p(i_1) = p(i_2) = -1, \\ 1 & \text{else} \end{cases} \quad (9)$$

with the two incoming indices i_1, i_2 and the two outgoing indices j_1, j_2 .

In our graphical notation, we place a SWAP gate on every line crossing. SWAP gates are pictured as diamonds in our notation as used in Figs. 1 and 2. We can move the lines freely and create new intersections with SWAP gates due to S being self-inverse. Because of Eq. (7), we can also move lines across nodes. When we want to apply nearest-neighbor gates, we can move the physical links to be next to each other. This introduces new SWAP gates in the network but allows the operators to be applied locally.

The SWAP gates were used in Fig. 2 to write a bra state in an alternative manner. Originally, the bra state can be obtained by reflecting the ket state on a line in the vertical direction and subsequently conjugating all tensors. This is shown as the left side of Fig. 2(a). The right side arranges the ket state in the same way as the bra state. This can be realized by appropriately moving the tensors and lines, and introducing new SWAP gates. Some of the SWAP gates are included in the definition of the starred tensors according to Fig. 2(b).

In this manner we never have to contract the SWAP gate with a tensor explicitly. Instead, we can always find a way to write the tensor network so that the gate is contracted with two indices of the same tensor. This way the SWAP gate can be applied element wise to the tensor. More details can be found in Appendix A. The computational cost scales only with the number of elements of the corresponding tensor. Since this is negligible compared to the most demanding steps we describe later in Sec. IIE, the fermionic algorithm has the same leading costs as the bosonic one.




FIG. 3. Splitting tensors into triads *triad* and the environments *env* for an update in the x direction. Symbology: small orange circles on indices, singular values; diamond, swap gate for fermionic PEPS.

2. Parity index

The previously constructed fermionic PEPS requires that all tensors have an overall even parity due to Eq. (7). This way, the tensor network can only represent states with even parity. To include the odd-parity sector as well, one introduces an additional parity index to one of the tensors. This index can take two values corresponding to even or odd parity. We represent it by dotted lines (see, e.g., Fig. 1). When calculating expectation values, the parity index of a bra state gets contracted with the corresponding one in the ket state.

We can also crop the parity index to only have one value with either even or odd parity. This way we simulate the even- and odd-parity sectors individually. For the Hubbard model, the state of half-filling is of even parity. By choosing an odd parity, we can find the state with an additional electron (or hole).

D. Truncation and expectation values

To apply the imaginary-time evolution described in Sec. II B, the nearest-neighbor operators $e^{-H_i \delta t}$ have to be applied to the PEPS sequentially. Since every application of a local operator increases the bond dimension on the corresponding link, a truncation has to be done in order to keep the bond dimension small. We use two different schemes for the truncation: simple update [8,28–30] and full update [8,30–32]. We limit ourselves to open boundary conditions (OBC) here, though we note that the simple update scheme can easily be applied to periodic boundary conditions as well. All results presented in this paper are obtained with simple update where not stated otherwise.

1. Simple update

Simple update is a truncation method that only takes local properties into account. We use the simplified update, where triads (reduced tensors) are split from the PEPS tensors to reduce the numerical costs [8,30,32].

In a first step, each tensor on the updated link is split into a triad and an environment tensor. This is depicted in Fig. 3 for the x direction, where the y direction gets split similarly.




FIG. 4. Simple update truncation.

The link to be updated gets combined with the physical index, while all other indices get combined to a second index. The resulting matrix is split in a QR decomposition into an orthogonal matrix Q and an upper triangular matrix R . The triangular matrix carries the physical index and the link to be updated, while the orthogonal matrix carries all other indices. The newly emerged index connecting the two matrices of the QR decomposition is denoted with wiggly lines in our graphical notation. We limit ourselves to updating the triads only as this is numerically cheaper. For fermionic PEPS we apply the SWAP gates to the initial tensors to put the physical indices next to each other. The two-site operator can then be directly applied to these adjacent indices.

The update is performed as depicted in Fig. 4. We apply the local gate to the physical indices of the triads. The resulting network gets contracted and is then split by a singular value decomposition (SVD) into two tensors U and V connected by a singular value matrix S . The singular values are truncated in order to keep the bond dimension bound by D . Finally, the square roots of the singular values can be multiplied to the left- and right-unitary tensors U and V to form the new triads. We keep the singular values explicitly on the links, however, as shown in Figs. 3 and 5 as small orange circles on the links. To increase the precision, the singular values of the surrounding links of tensors are included in the initial QR decomposition to calculate the triads and the environment (see Fig. 3). This way, more information about the rest of the tensor network is included during the update.

Finally, we update the PEPS tensors by recombining the new triads and the old environment tensors. This is shown in Fig. 5.




FIG. 5. Combining triads U , V and the environments *env* to the updated PEPS tensors $\tilde{L}$ and $\tilde{R}$.




FIG. 6. Overlap of the old and the new boundary MPS. The upper layer corresponds to the old boundary MPS, the middle layer represents bra and ket PEPS tensors, and the lower layer is the conjugate of the new boundary MPS.

For fermionic PEPS we split the tensor on the left side of Fig. 4 into two parts. Each one contributes to only one parity sector on the link to be updated. Then, two individual SVDs are performed on these parts. The truncation combines and orders the singular values and keeps the D largest ones. This allows to change the splitting between even and odd parity due to their local contributions to the tensor network.

With this update procedure we also change the norm of the state. In order to keep the norm in a reasonable order, we normalized the new singular values:

$$S \rightarrow \frac{S}{\lambda}, \quad \lambda = \sqrt{\sum_i S_i^2}. \quad (10)$$

The renormalization factors λ can also be used to estimate the energy (see Sec. III C).

2. Full update, boundary matrix product state, and expectation values

Full update algorithms take the whole tensor network into account for the truncation. Contracting the complete network scales exponentially in the size of one dimension of the system. To avoid this, an approximate contraction has to be used. We use the boundary matrix product state (boundary MPS) method for two-dimensional systems with OBC [21,30].

A boundary MPS is a string of tensors each connected to its neighbors. It represents an approximation of the contraction of the bra- and ket-layer of all tensors in the network from the boundary up to a certain x or y value. The upper layer of Fig. 6 represents a boundary MPS in the x direction. Each tensor is connected to its neighbor by an index (thick lines) of size up to χ , the boundary MPS dimension. Additionally, each tensor carries two indices which represent the external indices of the approximated network. Each is of size up to D .

Initially, the boundary MPS is set to a trivial product state of length L_x with all tensors set to rank 0, i.e., scalar. To go from one row to the next, we initialize the boundary MPS including one more row with random numbers. Then, a sweep is done through all tensors of the boundary MPS to optimize one tensor at a time, while the others are kept fixed. If $|b\rangle$ is the contraction of the old boundary MPS with the PEPS tensors to be included in the next layer, and $|b'\rangle$ is the new boundary MPS, we minimize

$$\| |b\rangle - |b'\rangle \|^2 = \langle b|b\rangle + \langle b'|b'\rangle - \langle b|b'\rangle - \langle b'|b\rangle. \quad (11)$$




FIG. 7. Calculating the expectation value of a two-site operator $\hat{o}$ in the x direction with a fermionic tensor network using boundary MPS. Wiggly lines denote a splitting of triads and environment tensors according to Fig. 5 but without singular values S . Contraction of all tensors outside the dashed-dotted rectangle defines the environment tensor N_{red} .

The overlap $\langle b|b'\rangle$ is shown in Fig. 6. The optimization can be done using alternating least-squares (ALS) optimization [30,32,33].

This procedure is done from all boundaries and allows us to calculate expectation values as depicted in Fig. 7. In the update step, the same surrounding can be used. We split the nodes we want to update into triads and the rest as described in Sec. II D 1. The whole network except for the triads is then contracted to form an environment tensor N_{red} as defined in Fig. 7.

For the truncation step, we make N_{red} Hermitian and semi-positive before we gauge it for better conditioning of the following steps. This preparation is explained in detail in [32]. For the update itself there are two common procedures. The first is called the second renormalization group (SRG) [34]. The idea here is to generalize simple update and do a singular value decomposition that includes the environment N_{red} . The truncation is done by dropping the smallest singular values of this decomposition. Since every singular value corresponds to either even or odd parity on the link to be updated, the splitting between the two parity sectors is done dynamically according to the sizes of the singular values in the sectors. The second update procedure optimizes the sites iteratively while keeping the remaining sites fixed, using the same ALS optimization techniques as before for the boundary MPS [30,32,33]. We applied this algorithm to fermionic PEPS, so we also had to sweep through the parity sectors. The details of the optimization can be found in Appendix B.

Despite its benefit of being able to dynamically track and split the parity sectors, we found that truncation errors of the SRG update process were rather large for the systems we tested and this precluded its use as a standalone updating process. On the other hand, the ALS procedure provided much smaller truncation errors, but the parity sectors have to be kept at a constant size during the updating process. This prevented us from dividing the parity sectors on a link according to their contribution to the state. Our solution to addressing the shortfall of each update procedure was to combine both procedures during the updating step, as follows: We first used SRG to make an initial guess for the new triads and to choose the size of the different parity sectors. The relatively large truncation errors at this point were subsequently reduced by application of the ALS updating procedure with a fixed size of the sectors. By combining both methods, we were able to achieve our

desired accuracy goals with full accounting of both parity sectors.

After each update step we divided the new state $|\Psi'\rangle$ by $\lambda = \sqrt{\langle\Psi'|\Psi'\rangle/\langle\Psi|\Psi\rangle}$ to normalize it.

E. Resource scaling

Now that we know which operations have to be performed, we can analyze the required resources in terms of runtime and memory. In both cases, the leading contribution is made by the calculation of the environment N_{red} as shown in Fig. 7. The overall runtime T_{CPU} then asymptotically scales as

$$T_{\text{CPU}}^{\square} \propto N_{\text{in}}(2\chi^3 D^4 d^2 + 2\chi^2 D^6 d^2 + \chi^2 D^4 d^4) \quad (12)$$

for a 2D square lattice where $N_{\text{in}} = (L_x - 2)(L_y - 2)$ denotes the number of inner lattice points. We do not need the full number of points here because the contraction at the boundary proceeds much faster due to the lower number of links on the tensors.

In the case of the honeycomb lattice some of the links are removed. This can be interpreted (and implemented) as setting $D = 1$ for these links. The runtime reduces to

$$T_{\text{CPU}}^{****} \propto N_{\text{in}}(2\chi^3 D^4 d^2 + 2\chi^2 D^5 d^2 + \chi^2 D^4 d^4). \quad (13)$$

It is common to choose $\chi \propto D^2$ which leads to the well-known and feared scaling in $O(D^{10})$. Our findings in Sec. III A, however, suggest that choosing $\chi \propto D$ (usually $\chi = 2D$ or $\chi = 3D$) already yields very good results in many cases. In particular, the error on observables introduced by the choice of χ is much smaller than that introduced by the choice of D . With $\chi \propto D$ the overall scaling of our algorithm reduces to $O(D^8)$ for the square and $O(D^7)$ for the honeycomb lattice.

For practical considerations such as job submissions, an understanding of the required memory M is often even more important than the scaling of the runtime. We found that we can predict the memory usage very accurately [up to $O(100 \text{ MB})$] with the formula

$$M = 700 \text{ MB} + 2\chi^2 D^4 d^2 M_{\text{num}} + 2N\chi^2 D^2 d M_{\text{num}}, \quad (14)$$

where M_{num} is the memory required for a single number, in our case $M_{\text{num}} = 16 \text{ B}$ for complex doubles. Notice that the constant offset in Eq. (14) is system dependent and in our case represents the approximate memory usage of MATLAB [35] when idle.

Our simulations were single threaded and accomplished within 48 hours. They were performed on MATLAB with heavy reliance on NCON [36].

F. Hubbard model

The Hamiltonian of the Hubbard model can be decomposed into sums of operators acting on a pair of neighboring sites only. It reads as

$$H = -\kappa \sum_{\langle x,y \rangle, s} c_{x,s}^{\dagger} c_{y,s} + \frac{U}{2} \sum_x q_x^2 + \mu \sum_{x,s} \left(c_{x,s}^{\dagger} c_{x,s} - \frac{1}{2} \right) \quad (15)$$

with $\langle x, y \rangle$ denoting nearest neighbors, whereas $s \in \{\uparrow, \downarrow\}$ denotes the two possibilities for the spin of an electron. $c_{x,s}^{\dagger}$

$(c_{x,s})$ is a creation (annihilation) operator of an electron at position x with spin s and follows the usual anticommutation relations

$$\{c_i^{\dagger}, c_j^{\dagger}\} = \{c_i, c_j\} = 0, \quad \{c_i, c_j^{\dagger}\} = \delta_{ij}. \quad (16)$$

The onsite coupling U gives the interaction strength of two electrons with combined charge $q_x = c_{x,\uparrow}^{\dagger} c_{x,\uparrow} + c_{x,\downarrow}^{\dagger} c_{x,\downarrow} - 1$ on the same lattice point. The hopping parameter is denoted by κ and μ is the chemical potential. The form of Eq. (15) is such that at $\mu = 0$ the system is at half-filling.

The Hubbard model is studied on a honeycomb lattice. The connectivity of nearest neighbors $\langle x, y \rangle$ is the same as the connectivity of the tensor network, as can be seen in Fig. 1. For larger systems in the x (y) direction, the figure is extended to the right (bottom) in our convention. The parity link is always attached to the tensor in the left bottom corner.

We rewrite the Hamiltonian Eq. (15) purely as a sum of two-site terms h_{xy} as follows:

$$H = \sum_{\langle x,y \rangle} h_{xy}, \quad (17)$$

$$h_{xy} = -\kappa \sum_s (c_{x,s}^{\dagger} c_{y,s} + c_{y,s}^{\dagger} c_{x,s}) + \frac{U}{2} \left(\frac{1}{n_x} q_x^2 + \frac{1}{n_y} q_y^2 \right) + \mu \sum_s \left[\frac{1}{n_x} \left(c_{x,s}^{\dagger} c_{x,s} - \frac{1}{2} \right) + \frac{1}{n_y} \left(c_{y,s}^{\dagger} c_{y,s} - \frac{1}{2} \right) \right]. \quad (18)$$

Here n_x is the number of nearest neighbors of site x . h_{xy} is chosen Hermitian and therefore every link $\langle x, y \rangle$ is counted only once, i.e., $\langle y, x \rangle$ is ignored. For the imaginary-time evolution we then use the Hamiltonian gate $e^{-h_{xy}\delta t}$.

1. Two-site basis

We now define our basis for the two-site states. For this we fix the order of creation operators acting on the empty state $|0\rangle$ as $c_{x,\uparrow}^{\dagger}, c_{x,\downarrow}^{\dagger}, c_{y,\uparrow}^{\dagger}, c_{y,\downarrow}^{\dagger}$ and receive a $(4 \otimes 4 = 16)$ -dimensional basis, where the one-site basis consisting of $|0\rangle, |\uparrow\downarrow\rangle, |\uparrow\rangle$, and $|\downarrow\rangle$ is $d \equiv 4$ dimensional. We write

$$|k\rangle := |k_x\rangle_x \otimes |k_y\rangle_y, \quad (19)$$

$$k = 4k_x + k_y \quad (20)$$

and the $k_{x,y} \in \{0, \dots, 3\}$ is the respective state index, so for example $|7\rangle = |1\rangle_x \otimes |3\rangle_y \equiv |\uparrow\downarrow\rangle_x \otimes |\downarrow\rangle_y$.

In this basis we can rewrite the gate operator h_{xy} using the defining properties

$$\langle k_2 | c_{x,\uparrow}^{\dagger} c_{y,\uparrow} | k_1 \rangle = \begin{cases} 1 & (k_2, k_1) \in \{(8, 2), (11, 1)\}, \\ -1 & (k_2, k_1) \in \{(4, 14), (7, 13)\}, \\ 0 & \text{else,} \end{cases} \quad (21)$$

$$\langle k_2 | c_{x,\downarrow}^{\dagger} c_{y,\downarrow} | k_1 \rangle = \begin{cases} 1 & (k_2, k_1) \in \{(12, 3), (4, 11)\}, \\ -1 & (k_2, k_1) \in \{(14, 1), (6, 9)\}, \\ 0 & \text{else} \end{cases} \quad (22)$$

and

$$q_x^2 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \otimes \mathbb{1}_4, \quad (23)$$

$$c_{x,\uparrow}^\dagger c_{x,\uparrow} + c_{x,\downarrow}^\dagger c_{x,\downarrow} - 1 = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \otimes \mathbb{1}_4. \quad (24)$$

2. Operators

Any local operator, i.e., involving only onsite and nearest-neighbor interactions, can be expressed in terms of gates in the same way as the Hamiltonian in Eq. (17). We derive this form explicitly for the magnetization operator because it becomes important for breaking degeneracies later on. The operator reads as

$$M = \sum_x (c_{x,\uparrow}^\dagger c_{x,\uparrow} - c_{x,\downarrow}^\dagger c_{x,\downarrow}) \quad (25)$$

$$\begin{aligned} &= \sum_{(x,y)} \left(\frac{1}{n_x} (c_{x,\uparrow}^\dagger c_{x,\uparrow} - c_{x,\downarrow}^\dagger c_{x,\downarrow}) \right. \\ &\quad \left. + \frac{1}{n_y} (c_{y,\uparrow}^\dagger c_{y,\uparrow} - c_{y,\downarrow}^\dagger c_{y,\downarrow}) \right) \\ &=: \sum_{(x,y)} m_{xy} \end{aligned} \quad (26)$$

with the matrix form

$$\begin{aligned} m_{xy} &= \left[\frac{1}{n_x} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \otimes \mathbb{1}_4 \right. \\ &\quad \left. + \frac{1}{n_y} \mathbb{1}_4 \otimes \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \right]. \end{aligned} \quad (27)$$

The integers n_x and n_y denote the number of nearest neighbors at sites x and y .

The particle number

$$n = \sum_x (c_{x,\uparrow}^\dagger c_{x,\uparrow} + c_{x,\downarrow}^\dagger c_{x,\downarrow}) \quad (28)$$

is another important operator as it directly couples to the chemical potential. Its matrix form follows immediately from Eq. (24).

III. IMPROVEMENTS

A. Boundary MPS bond dimension dependence

Since our observables consist of nearest-neighbor operators only, one can use the boundary MPS environment for each set of nearest neighbors and calculate the corresponding expectation value as explained in Sec. II D and Fig. 7. The sum of these terms leads to the desired global observables.

In the literature a typical value of $\chi \geq D^2$ is standard [8,21,29,31,32,37]. This gives an exact representation of the boundaries itself and a good approximation of the rest of the




FIG. 8. Dependence of different operators $\langle O \rangle$ on the boundary MPS bond dimension χ for the ground state in the odd-parity sector. Shown are the relative errors compared to results at $\chi = 100$. Honeycomb Hubbard model with $D = 12$, $L = 12 \times 6$, $\kappa = 1$, $U = 2$, $\mu = 0.1$, $B = 0.01$. See [38] for numerical data.

network. We considered the dependence of expectation values on the boundary MPS bond dimension more carefully and found that a much smaller χ leads to a good precision already.

For example, Fig. 8 shows the relative error of different observables as a function of χ . In this case, the ground state of the odd-parity sector of the Hubbard model was obtained with simple update and a bond dimension of $D = 12$. Expectation values were then calculated with different boundary MPS dimensions χ and compared to the result using $\chi = 100$. We found that for $\chi \geq 16$ all errors are on a subpercent level. This is much smaller than the usual $D^2 = 144$. Thus, if a subpercent precision is sufficient, as it was in our studies, the boundary MPS bond dimension can be decreased significantly. We found similar results for different parameters U and μ , bond dimensions, and both parity sectors of the Hubbard model.

Let us stress, however, that these results only apply to the boundary MPS contraction after a time evolution with simple update. They do not apply to full update. On the contrary, for full update to be stable we found that in some cases $\chi \gtrsim D^2$ is indeed required.

Obviously, choosing a smaller χ improves the computational costs drastically. This in turn can be used to increase D which then improves the overall precision. Note that we need to normalize our state or calculate the norm explicitly for expectation values:

$$\langle \hat{O} \rangle = \frac{\langle \Psi | \hat{O} | \Psi \rangle}{\langle \Psi | \Psi \rangle}. \quad (29)$$

Instead of normalizing the state once, we found it helpful to calculate the norm $\langle \Psi | \Psi \rangle_k$ for each nearest-neighbor pair k

by calculating it the same way as the expectation values, but without inserting any operators. This way, the local expectation values can be divided by the norm with respect to the same environment N_{red} and therefore the same truncation of the tensor network. This increases the precision of expectation values. Moreover, the standard deviation ΔI of the local norm evaluations can be used as an estimator for the error introduced by the boundary MPS truncation:

$$\Delta I = \frac{\sqrt{\sum_k (\langle \Psi | \Psi \rangle_k - \overline{\langle \Psi | \Psi \rangle})^2}}{\overline{\langle \Psi | \Psi \rangle} \sqrt{|k| - 1}}. \quad (30)$$

The term $|k|$ denotes the number of nearest-neighbor pairs in the system and $\overline{\langle \Psi | \Psi \rangle} = \frac{\sum_k \langle \Psi | \Psi \rangle_k}{|k|}$ is the mean value of the norm. The error estimation by ΔI can be seen in Fig. 8 as red squares. In all cases we studied, ΔI was of the order of the relative errors of any observable considered. This allowed us to estimate the effects of the boundary MPS truncation on expectation values without needing to do a full scan in χ .

We observed that a good precision can be reached with $\chi = 3D$. The error introduced by a finite D is then larger than that caused by a finite χ . A similar behavior was observed for infinite PEPS in [39]. We checked that our results do not crucially depend on χ by repeating most of our simulations with $\chi = 2D$.

B. Convergence

The imaginary-time evolution as described in Eqs. (1)–(5) has to be tuned carefully. It is crucial to choose an appropriate value for δt at every step. To understand the details better, we first have to differentiate two kinds of errors. First, we have the pollution of the ground state² by excited states. This error arises from Eq. (1), decays exponentially for any given δt , and we therefore call it *exponential error*. Second, the discretization into a scheme of order $r - 1$ introduces an error decreasing as δt^r [see Eq. (4)]. This error remains for any given δt even at $t \rightarrow \infty$. We designate this as the *Trotter error*.

A well-tuned convergence scheme should use as few steps as possible to achieve the smallest total error. Naively, this can be achieved by choosing several, successively decreasing, values of δt and performing the complete imaginary-time evolution for every one of them. This, however, turns out to be highly inefficient. Instead, we first perform the imaginary-time evolution with a coarse step size and then proceed from the resulting state with a smaller value of δt . As the ground state energies of both discretizations are (asymptotically) close to each other, we only have a small number of steps to perform with the finer step size. Once the new discretization shows no significant exponential error any more, the step size is reduced again and the procedure is iterated.

Thus, the overall convergence behavior resembles a stair as can be seen in Fig. 9. If the “steps” of this stair have a significant slope at the time a smaller δt is chosen, this indicates a remnant exponential error and thus a loss of accuracy

²Here we mean the ground state of the Trotterized time-evolution operator.




FIG. 9. Relative deviations of the energies during the imaginary-time evolution with $\kappa = 1$, $U = 4$, and $\mu = B = 0$ on the 2×2 lattice with $D = 6$. (a) Too rapid reduction of δt . (b) Efficient reduction of δt . (c) Too slow reduction of δt . See [38] for numerical data. In all three cases exactly 50 steps have been performed. Energies calculated using boundary MPS (bMPS) and the direct estimator derived in Sec. III C.

[see Fig. 9(a)]. Therefore, δt has to be decreased slower. If on the other hand there are long plateaus or the energy difference between adjacent “steps” is very small, no accuracy is lost, but the algorithms efficiency could be increased by a more rapid decrease of δt [see Fig. 9(c)].

1. Extrapolation

For energy values $E(\delta t)$ with pure (or at least strongly dominating) Trotter error we make the ansatz

$$E(\delta t) = E_0 + \gamma \delta t^r + O(\delta t^{r+1}), \quad (31)$$

where E_0 is the exact ground-state energy and γ is some coefficient independent of δt . We then obtain an extrapolated E_0 by choosing two values of δt , preferably the smallest two available, and calculate

$$E_0 = \left(\frac{1}{\delta t_1^r} - \frac{1}{\delta t_2^r} \right)^{-1} \left(\frac{E(\delta t_1)}{\delta t_1^r} - \frac{E(\delta t_2)}{\delta t_2^r} \right) + O(\delta t^{r+1}). \quad (32)$$

2. δt -reduction scheme

The energy including the exponential error can be approximated by

$$E(t) = E_0(1 + \alpha e^{-\Delta_{\text{tr}} t}), \quad (33)$$

where α is some coefficient independent of t (and at least asymptotically of δt) and Δ_{tr} is the current gap to the first excited state. This approximation is accurate up to exponentially suppressed higher-state contributions and therefore very precise at large enough times. At large t and small δt the gap approaches the difference in Trotter errors between the last δt and the current one. We therefore write

$$\Delta_{\text{tr}} \xrightarrow{t \rightarrow \infty, \delta t \rightarrow 0} \beta \Delta_{\text{ex}} \delta t^r \quad (34)$$

with another coefficient β constant in t and δt . The physical gap Δ_{ex} to the first excited state³ has been included to set a comparable energy scale.

We consider the local relative energy change

$$E'(t) := \frac{E(t + \delta t) - E(t)}{E(t)} \quad (35)$$

$$= \frac{\alpha e^{-\Delta_{\text{tr}} t} (e^{-\Delta_{\text{tr}} \delta t} - 1)}{1 + \alpha e^{-\Delta_{\text{tr}} t}} \quad (36)$$

$$= \alpha e^{-\Delta_{\text{tr}} t} [-\Delta_{\text{tr}} \delta t + O((\Delta_{\text{tr}} \delta t)^2)] + O(e^{-2\Delta_{\text{tr}} t}) \quad (37)$$

$$\rightarrow \underbrace{-\alpha \beta e^{-\Delta_{\text{tr}} t}}_{\varepsilon} \Delta_{\text{ex}} \delta t^{r+1} \quad (38)$$

and find that at any given discretization the exponential error can be considered small enough once some threshold

$$|\varepsilon| = \left| \frac{E'(t)}{\Delta_{\text{ex}}} \right| \delta t^{-(r+1)} < \varepsilon_{\text{tol}} \quad (39)$$

is reached. The target precision of the exponential error ε_{tol} therefore can be chosen independently of t and δt .

The remaining problem is that *a priori* the factors β and especially Δ_{ex} are not known. We observe that β is usually of the order of one and therefore does not greatly influence the performance of our algorithm. Δ_{ex} on the other hand can vary

over many orders of magnitude if the parameters of the model are changed and a good initial guess can drastically improve the performance of the algorithm.

As explained earlier, we reduce δt once the exponential error decreases below the threshold set by ε_{tol} . We chose to perform this reduction geometrically, i.e., $\delta t \mapsto \xi \delta t$ with some factor $\xi \in (0, 1)$. It is important to choose ξ carefully. If it is chosen too small, unnecessarily many steps are performed with the small step size where a few larger steps would have sufficed. On the other hand, too large values of ξ lead to a slow decrease of δt implying a slow continuous time convergence. We find that usually the results are best for $\xi \in [0.7, 0.9]$.

3. Convergence criteria

Our termination criterion for the evolution is the point where the exponential error is diminished as explained before and δt has dropped below some threshold. This threshold usually can be chosen surprisingly high because the extrapolation of the values without exponential error reduces the Trotter error to $O(\delta t^{r+1})$.

C. Energy estimator

The energy can be measured as described before by calculating the expectation values of the local terms of the Hamiltonian as in Fig. 7. This requires the calculation N_{red} , which is the leading cost of the whole algorithm, for every site. If we use a full update algorithm, N_{red} has to be calculated in each update step as well. Therefore, calculating the energy is as expensive as a step of the imaginary-time evolution. It can be even cheaper if parts of the previously calculated boundary MPS are reused.

The situation looks different, however, when we use simple update. Each update step scales at most as $O(D^4)$ in the bond dimension for the honeycomb lattice. This is much cheaper than the calculation of an expectation value. Hence, we used an estimator for the energy to avoid calculating N_{red} after each update step.

From Eq. (2) follows (up to discretization errors)

$$H = -\frac{1}{\delta t} \ln U(\delta t) \quad (40)$$

$$\Rightarrow E = -\frac{1}{\delta t} \frac{\langle \Psi | \ln U(\delta t) | \Psi \rangle}{\langle \Psi | \Psi \rangle}. \quad (41)$$

If $|\Psi\rangle$ is the ground state of the system, we can approximate this as

$$E \approx -\frac{1}{\delta t} \ln \sqrt{\frac{\langle \Psi | U(\delta t)^2 | \Psi \rangle}{\langle \Psi | \Psi \rangle}} \quad (42)$$

$$= -\frac{1}{\delta t} \ln \sqrt{\frac{\langle \Psi' | \Psi' \rangle}{\langle \Psi | \Psi \rangle}} = -\frac{1}{\delta t} \ln \prod \lambda. \quad (43)$$

In the last line we named the new state after an imaginary-time step $|\Psi'\rangle$. The renormalization factor λ corresponds to the change of the norm of a state when an imaginary-time evolution operator is applied locally [see Eq. (10)]. In the product $\prod \lambda$ we combine the individual renormalization factors from all two-site updates involved in a global update. With this

³Only the gap in the simulated sector is relevant here. If, e.g., only the even-parity sector is simulated, then Δ_{ex} denotes the gap between even ground state and first excited even state.




FIG. 10. Imaginary-time evolution with $\kappa = 1$, $U = 4$ on the 2×4 lattice with $D = 8$ and odd parity. Energies calculated using boundary MPS (bMPS) and the direct estimator derived in Sec. III C. (a) $\mu = B = 0$; (b) $\mu = B = 0.1$; see [38] for numerical data. As a reference we provide the exact imaginary-time evolution of the state vector obtained via full contraction of the initial PEPS.

procedure we get an estimator of the energy from the simple update algorithm without significant additional costs.

The estimator is only exact up to the Trotter error. It also assumes that we converged to the ground state. In some cases the estimator reproduces the energy quite reliably even if convergence or a small δt are not given. In other cases, the estimator was far off the real energy. See Figs. 9–11 for examples.

We found the estimator useful to check for convergence within a given Trotter step size δt . If the energy estimator does not change anymore after an imaginary-time step with a given step size, the state converged to the ground state of the Trotterized Hamiltonian. Therefore, we can reduce the step size. This convergence check is much cheaper than the calculation of an expectation value after each imaginary-time step.




FIG. 11. Energy during the imaginary-time evolution with $U = 1$ and $\kappa = \mu = B = 0$ on the 3×4 lattice with odd parity. The energies have been calculated using boundary MPS (bMPS) and the direct estimator derived in Sec. III C for simple update. Full update simulation with bMPS energy calculation. See [38] for numerical data.

D. Instabilities caused by degenerate states

Even if all the criteria described above are fulfilled (including a good estimation of the gap Δ_{ex}), the imaginary-time evolution is still not guaranteed to converge to the correct value. As a matter of fact, it might not converge at all. One such example is shown in Fig. 10(a), where the results first approach the correct values. At some point, however, the energy estimator does not decrease monotonously any more. It begins to rise and to fluctuate. Again, much later the results become numerically unstable and chaotic.

Fortunately, the onset of this problem has various indicators which we can track. Not only does the full contraction energy estimator fail to decrease further, in our implementation it also gains a significant imaginary part as compared to machine precision. Under normal circumstances, the imaginary part is on the order of 10^{-13} , but when the time evolution starts to fail it increases drastically. It becomes clear that we are facing numerical instabilities when taking into account that even the norm estimator obtains an imaginary part. One further, though less reliable, indicator is a steplike decrease of the simple update energy estimator. The step does not always occur and it sometimes stabilizes at a new level, but other times it continues decreasing without bounds.

We find that such instabilities occur only if the ground state of the simulated system is degenerate. It is plausible that degeneracies lead to numerical instabilities because the ambiguity of the unitary matrices in an SVD is dramatically increased once the principal values are no longer pairwise distinct. This is similar to simulations without gauge fixing which are known to be unstable. We find that breaking the degeneracy, if ever slightly, resolves the problem. In Fig. 10(a) the simulated system has a fourfold-degenerate ground state and we clearly observe the described instability. With all symmetries broken (see next section for details) and a non-degenerate ground state the evolution converges stably as can be seen in Fig. 10(b).

1. Degeneracies in the Hubbard model

The ground state of the Hubbard model at half-filling (i.e., $\mu = 0$) lies in the even-parity sector and is nondegenerate [14]. Therefore, simulations of the full space and the even-parity subspace do not feature the instabilities described above. One has to be careful nevertheless because influences like a strong chemical potential can reintroduce degeneracies. In the case of large μ the ground state falls into a sector with nonzero⁴ particle number which might be degenerate.

The ground state of the odd-parity sector, on the other hand, *a priori* comes with a fourfold degeneracy. It can have one spin-up or -down particle less or more than the ground state, all leading to the same energy at $\mu = 0$ with no external magnetic field B . This degeneracy is easily broken by introducing a small but nonzero chemical potential and a magnetic field of similar magnitude. The magnitude has to be chosen carefully as it has to be large enough to ensure numerical stability but small enough to keep the ground state in the desired sector. In practice, $\mu \sim 10^{-2}\kappa$ usually yields good results and the quality of the simulation is not very sensitive to the exact value of μ .

The ground-state energy E_0 without chemical potential and magnetic field can be easily regained after a simulation from the measured one:

$$E_0(\mu, B) = E_0 - |n\mu| - |MB|. \quad (44)$$

We exploit the fact here that particle number n and magnetization M are both good quantum numbers of the Hubbard model, i.e., the corresponding operators commute with the Hamiltonian. We use the absolute values in Eq. (44) for clarity as the ground state of the broken system always has lower energy than the ground state of the unbroken one.

Finally, we point out that the lattice geometry itself can exhibit an additional, more subtle, degeneracy that is not as easily broken and regained after the simulation. This degeneracy comes from a high spatial symmetry of the lattice or, put differently, is a degeneracy in momentum space. It occurs only if the symmetry group of the lattice has at least one irreducible representation of dimension larger than one. In this case, the degeneracy is protected by symmetry and cannot be broken by any translationally invariant Hamiltonian. Any ring with an even number N of sites exhibits this feature as it falls into the dihedral symmetry group D_N . In our case of rectangular pieces of the hexagonal lattice with OBC it turns out that the 6-ring is the only lattice with such a high symmetry. All the other lattices can only have reflection symmetries and thus only one-dimensional irreducible representations. We can therefore easily avoid this degeneracy by choosing the geometry accordingly from here on.

E. Instabilities caused by the single-particle gap

A different problem arises in the case of a very large single-particle gap $\Delta_{\text{sp}} \gg \Delta_{\text{ex}}$. The single-particle gap is the

difference between even- and odd-parity ground-state energies E_0^e and E_0^o , respectively. It turns out that in this regime the simulation in the even-parity sector works fine, but the odd-parity simulation becomes unreliable. During the time evolution the energy in the odd-parity sector can drop to that of the even-parity sector, for example. To demonstrate this type of instability explicitly, we simulate in the regime where $\kappa = 0$, thereby enforcing $\Delta_{\text{sp}} \gg \Delta_{\text{ex}}$.

Such a case is shown in Fig. 11 where an odd-parity simulation has been attempted. The energy estimator from the change of the norm during the simple update goes down to $E_0^e = 0$ immediately. The energy estimator from the boundary MPS contraction originally converges correctly against $E_0^o = \frac{1}{2}U$, but then jumps down to E_0^e after some time $t \approx 35$. Only the simulation using full update remains stably at the correct energy.

The chosen bond dimension $D = 2$ suffices to solve the problem at hand exactly because with no hopping the Hamiltonian decomposes into its local contributions. On every lattice point, there is a twofold-degenerate ground state (one particle up or down) with energy 0 and a twofold-degenerate excited state (none or two particles) with energy $U/2$. That means that only two independent dimensions have to be considered. Thus, the instability cannot be alleviated with larger D .

1. Tensor structure

To understand how this can happen, let us revisit the structure we imposed onto the tensors by the condition that only parity-conserving entries can be nonzero. We consider the simplest possible example, namely, a two-site lattice with OBC and $d = D = 2$. Then, both tensors can be written as 2×2 matrices. We add a parity index to the left tensor and fix it to be odd. Equation (7) now guarantees that the full tensor network takes the form

$$M_1 M_2 := \begin{pmatrix} 0 & a_1 \\ b_1 & 0 \end{pmatrix} \begin{pmatrix} a_2 & 0 \\ 0 & b_2 \end{pmatrix} \quad (45)$$

$$= \begin{pmatrix} 0 & a_1 b_2 \\ a_2 b_1 & 0 \end{pmatrix}, \quad (46)$$

which is equivalent to the state vector $v := (0, a_1 b_2, a_2 b_1, 0)$ encoding the two-dimensional odd subspace of the ($d^2 = 4$)-dimensional state space. Next, we apply a time-evolution operator U to both sites independently. Again to keep things simple, we choose $U = \text{diag}(1, \epsilon)$ where we are interested in the case of $\epsilon \ll 1$. The evolved network reads as

$$(U M_1)(M_2 U) = \begin{pmatrix} 0 & a_1 \\ \epsilon b_1 & 0 \end{pmatrix} \begin{pmatrix} a_2 & 0 \\ 0 & \epsilon b_2 \end{pmatrix} \quad (47)$$

$$= \begin{pmatrix} 0 & \epsilon a_1 b_2 \\ \epsilon a_2 b_1 & 0 \end{pmatrix}. \quad (48)$$

This means that ultimately the eigenstate v is multiplied by ϵ . The tensors $M_{1,2}$, however, are not “eigentensors” of the time evolution.

2. Suppression of a subspace

Equation (47) shows why the application of this energy estimator of Sec. III C using the local norm change is prob-

⁴In our convention, half-filling has exactly zero particles though of course the electron number has to equal the number of lattice sites. It is a simple shift by a constant.




FIG. 12. Energies of the 3×4 hexagonal lattice with $\kappa = 1$ and $\mu = B = 0$ at different values of the coupling U . (a) Even parity. (b) Odd parity. Duplicate points correspond to $\chi = 2D$ and $\chi = 3D$. The legend in the left plot applies to both. See [38] for numerical data.

lematic. Repeated applications of U lead to a suppression of the b terms. Once they are much smaller than the a terms, the matrix norm no longer changes significantly with further applications of U . The energy estimator thus becomes zero, or more generally E_0^e .

With this in hand, we can also estimate when the jump, i.e., the more important problem, will occur. We expect that the odd-parity sector will not be represented sufficiently well any more once the b terms reach the order of machine precision relative to the a terms. We call the threshold below which stability cannot be guaranteed any more $\varepsilon_{\text{stable}}$. This leads to the stability condition

$$e^{-\Delta_{\text{sp}} t} > \varepsilon_{\text{stable}} \quad (49)$$

$$\Rightarrow t < \frac{\ln(1/\varepsilon_{\text{stable}})}{\Delta_{\text{sp}}}. \quad (50)$$

In the example of Fig. 11 we find $\varepsilon_{\text{stable}} \approx 10^{-8}$.




FIG. 13. Smallest bond dimension D for which the simulation of the 3×4 hexagonal lattice with $\kappa = 1$ and $\mu = B = 0$ in the odd-parity sector becomes unstable against the ratio of single-particle and excitation gaps. Values of U as in Fig. 12 resulting in different ratios $\Delta_{\text{sp}}/\Delta_{\text{ex}}$. See [38] for numerical data.

One would expect $\varepsilon_{\text{stable}}$ to depend on D because more numerical errors can accumulate with a larger number of entries. This is not always the case, however, as can be clearly seen in Fig. 11. Unfortunately, this property is not generalizable and there are cases where larger D can worsen the problem. Such examples are visualized in Figs. 12–14 and explained in detail in Sec. IV A.

As the stability condition takes the same form as the convergence condition for the exponential error due to the gap Δ_{ex} , we can easily extract the ratio between the time t_{stable} for which a stable evolution is possible and the time t_{conv} needed to reach the precision $\varepsilon_{\text{conv}}$, giving

$$\frac{t_{\text{stable}}}{t_{\text{conv}}} = \frac{\Delta_{\text{ex}} \ln \varepsilon_{\text{stable}}}{\Delta_{\text{sp}} \ln \varepsilon_{\text{conv}}}. \quad (51)$$

The ratio of the logarithms can vary depending on the desired precision and the dimensionality of the system, but in all realistic scenarios it is of order 1 or slightly larger. Therefore,




FIG. 14. Standard deviation of the norm ΔI [see (30)] and deviations of the magnetization M [see (26)] and the particle number n [see (28)] from the exact value. 3×4 hexagonal lattice with $\kappa = 1$, $U = 3$, and $\mu = B = 0$ for different bond dimensions D using $\chi = 3D$. See [38] for numerical data.

the relevant factor is the ratio between the different gaps. One finds that a reliable simulation is possible if and only if $\Delta_{\text{ex}} \gtrsim \Delta_{\text{sp}}$.

Note that Eq. (50) provides a lower bound for t_{stable} , i.e., the time in which stability is guaranteed. It is nevertheless possible that simulations proceed stably for a much longer time. We observe that this is in fact usually the case. A possible explanation is that interactions involving more than one site average out the effect of suppressed subspaces on a larger part of the lattice and thus mitigates the problem. The example above has no such interactions and hence accumulates the factors locally.

3. Adding the even-parity sector

The example above Eq. (46) assumes an odd external parity index at M_1 . In the more general case we would allow this external index to take both values and therefore obtain two copies of the matrix M_1^o and M_1^e where the odd version M_1^o corresponds to M_1 in Eq. (46). In this case, an imaginary-time step evolves the network as

$$M_1^o M_2 \sim \begin{pmatrix} 0 & 1 \\ \epsilon & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & \epsilon \end{pmatrix} \quad (52)$$

$$= \begin{pmatrix} 0 & \epsilon \\ \epsilon & 0 \end{pmatrix} \quad (53)$$

in the odd-parity part and as

$$M_1^e M_2 \sim \begin{pmatrix} 1 & 0 \\ 0 & \epsilon \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & \epsilon \end{pmatrix} \quad (54)$$

$$= \begin{pmatrix} 1 & 0 \\ 0 & \epsilon^2 \end{pmatrix} \quad (55)$$

in the even-parity sector.

4. Full update as a possible fix

So far we are not aware of any solution to this problem as long as simple update is used. Full update, on the other hand, provides a solution as can be seen in Fig. 11, be it a very costly one. In the future hybrid solutions might be considered performing several steps with simple update and then a stabilizing step with full update. We believe that full update does not suffer instabilities from this problem because the full contraction of the network and the following normalization prevents the scale separation of individual blocks as in Eq. (47). More research, however, is required and ongoing with respect to full update's stability and possible alternative fixes.

IV. RESULTS

We now apply the algorithms explained above to specific test models. We provide most of our results for the 3×4 honeycomb lattice with OBC because this is the largest system for which results from exact diagonalization are feasible for a comparison. Some tests on larger lattices as a proof of principle are presented as well.

We find that the common choice of $\chi = D^2$ is much larger than required (see Sec. III A). A good convergence is observed

already at $\chi = 3D$. From now on all our simulations are performed with both $\chi = 2D$ and $3D$. Usually, the deviation between these two points is negligible compared to the differences between results at different D . We use simple update in the imaginary-time steps.

A. Ground state and first excited state of the Hubbard model

We first demonstrate the capabilities and limitations of our algorithm with a scan of different onsite couplings U which provides a broad range of single-particle gaps Δ_{sp} . The results for both even- and odd-parity sectors are depicted in Fig. 12. While the even-parity calculation yields reliably good results (with errors in the subpercent level for $D \geq 10$) for all chosen couplings, this is not the case in the odd-parity sector. There, the results are similarly good for $U \lesssim 2$ corresponding to the case of $\Delta_{\text{sp}} < \Delta_{\text{ex}}$. Above this threshold, however, the stability decreases. At $U = 2.5$ only simulations with bond dimensions $D \geq 20$ produce wrong results. The ever smaller bond dimensions at which the instabilities occur are plotted in Fig. 13.

Fortunately, the breakdown in the odd-parity simulations is easily identifiable. Not only the energy but also particle number n and magnetization M yield completely wrong results. Since the latter are known to be integer values in a nondegenerate ground state, they can be used for a simple consistency check. A large standard deviation of the norm ΔI provides an even more reliable indicator for untrustworthy results [see Eq. (30) and Fig. 8]. We visualize all three observables in Fig. 14 for a conditionally stable simulation. The region of $D \geq 12$, where the odd-parity sector features stability issues, is clearly distinct from the stable regions for all three indicators. There are two orders of magnitude separating the best unstable result from the worst stable one.

This observation implies that, even though stability problems can occur, we can still trust the results from odd-parity simulations as long as they have been tested for the stability indicators, in particular ΔI .

B. Finite chemical potential

To emphasize the efficacy of our tensor method when dealing with nonzero chemical potential μ , we compare our results to those obtained from hybrid (or Hamiltonian) Monte Carlo (HMC) for the 3×4 lattice with open boundaries. The nonzero value of μ induces a complex phase in the action, which we label θ for convenience.⁵ Standard HMC calculations must therefore resort to reweighting techniques when calculating some observable O :

$$\langle O \rangle = \frac{\langle O e^{i\theta} \rangle}{\langle e^{i\theta} \rangle}. \quad (56)$$

The expectation values on the right-hand side of the equation are evaluated with the real part of the action only. The reliability of such reweighting calculations depends strongly on the

⁵The phase θ is itself a function of the field configurations.




FIG. 15. The expectation value of the exponential of the complex phase θ of the action for the 3×4 honeycomb lattice as a function of chemical potential μ . Results are shown for integration along the real plane (black points) and integration pushed into the complex plane along the tangent plane (red squares) as described in the text. For $\mu = 0$ there is no sign problem and the expectation value is equal to 1 for both calculations. See [38] for numerical data.

fluctuations of $e^{i\theta}$ and this in turn influences its mean value $\langle e^{i\theta} \rangle$, or statistical power. In the limit of no complex phase (i.e., no sign problem), the statistical power goes to one and we recover the original HMC algorithm. On the other hand, in the limit of large fluctuations the statistical power goes to zero and the right-hand side of Eq. (56) becomes ill defined, precluding any reliable estimate of $\langle O \rangle$. In Fig. 15 we show the statistical power related to reweighting (black points) for this system as a function of μ . It is clear that the sign problem quickly becomes prohibitive for moderate $\mu \sim 0.3$ and larger. One can attempt to alleviate the sign problem by deforming the contour of integration into the complex plane [40–42]. The simplest contour deformation is that of a constant imaginary shift of the fields to a “tangent” plane that intersects the action’s main critical point [43,44]. Examples have been found where such a simple transformation of the fields greatly alleviates the sign problem [45]. We have done similar calculations with the honeycomb Hubbard system, as shown in Fig. 15 (red squares). Although there is slight improvement in the statistical power for $\mu \leq 0.3$, the sign problem still quickly overwhelms the calculation as μ increases.

For tensor network simulations this problem is completely absent. Figure 16 visualizes the energy obtained with different bond dimensions D in the even- and odd-parity sectors as well as the single-particle gap as the difference of these energies. The same system and the same range of chemical potentials as in Fig. 15 were used. The simulations have a lower precision in the regions where the ground state changes, i.e., near the kinks. This can be explained by the near degeneracy due to the small excitation gap Δ_{ex} . We also encounter severe problems in the odd-parity sector at $\mu = 0$ which is to be expected as the ground state in this sector is degenerate and we encounter the problem described in Sec. III D. Otherwise, the results are very precise, in particular, for a broad range of values $\mu > 0.3$ where the HMC completely fails.




FIG. 16. Energies of the 3×4 hexagonal lattice with $\kappa = 1$, $U = 2$, and $B = 0$ at different values of μ . (a) Even parity. (b) Odd parity. (c) Energy gap between even- and odd-parity sectors. Duplicate points correspond to $\chi = 2D$ and $\chi = 3D$. See [38] for numerical data.

C. Scalability with volume

The ability to simulate away from half-filling by itself is a great advantage of the tensor network ansatz over stochastic methods. Yet, the method is only beneficial if it demonstrates a feasible scaling in spatial volume. We know from Sec. II E that the runtime depends linearly on the lattice size and the volume’s influence on memory is of subleading order. It is,




FIG. 17. Relative deviations of the energies in the noninteracting limit with finite chemical potential ($\kappa = 1$, $U = 0$, $\mu = 0.5$, $B = 0$) from the exact values for different lattice sizes and bond dimensions D . Duplicate points correspond to $\chi = 2D$ and $\chi = 3D$. See [38] for numerical data.

however, not clear *a priori* that the same bond dimension D is enough to describe larger lattices with similar precision. Therefore, the implicit effect on runtime and memory might, in principle, be significantly larger. We find that this is not the case. As can be seen in Fig. 17, the relative error for any specific bond dimension D is virtually independent of the lattice size. This is a strong indicator that the area law is fulfilled for this system.

There is, however, another hidden cost of larger volumes. As the lattice approaches the thermodynamic limit, the excitation gap Δ_{ex} usually goes to zero. Thus, a longer imaginary-time evolution is required to reach convergence of the true ground state. In addition, it turns out that the time steps δt have to be chosen smaller from the very beginning for larger lattices. Too large time steps lead to numerical instabilities. The combination of these two constraints results

in a significantly increased number of steps required for the imaginary-time evolution.

We present the results of a 30×15 lattice⁶ simulation at nonzero chemical potential both in the noninteracting limit and at finite coupling $U = 2$ in Fig. 18. The first serves as a proof of principle where we can compare to results obtained from exact diagonalization of the hopping matrix. The latter simulation is the first of its kind and it predicts energies that could not have been obtained otherwise to date.

The results in Fig. 18 are depicted as a function of D^{-2} . We find empirically that this dependence approximates the convergence behavior very well. It should be possible to further improve the extrapolation by using the truncation error or similar quantities instead of the bond dimension [18,46].

Although we can calculate the energy well below sub-percentage precision in this way, we cannot give a reliable estimator of the single-particle gap Δ_{sp} because for this particular choice of parameters the gap is too small to be resolved, i.e., even- and odd-energy extrapolations are compatible within uncertainties. The results of the extrapolation in the noninteracting case are $E_{\text{even}} = -693.2(17)$ and $E_{\text{odd}} = -692.1(18)$ which in both sectors are fully compatible with the exact value $E_{\text{exact}} \approx -692.3$. The exact energy is the same for both parity sectors on our level of precision. In the case of $U = 2$ we find $E_{\text{even}} = -483.5(14)$ and $E_{\text{odd}} = -483.8(12)$. The fit has been done using the Levenberg-Marquardt algorithm [47,48]. Reference [49] provides a comprehensible explanation of the uncertainty estimation.

V. CONCLUSION

We presented an algorithm for direct simulations of fermionic even- and odd-parity sectors and demonstrate its efficacy using the example of the Hubbard model. For this, we use fermionic PEPS with an additional parity link confined to

⁶This corresponds, up to boundary conditions, to the bipartite 15×15 honeycomb lattice with two sites per unit cell.




FIG. 18. Energies with finite chemical potential ($\kappa = 1$, $\mu = 0.5$, $B = 0$) for the 30×15 lattice against the inverse squared bond dimension. (a) $U = 0$; (b) $U = 2$; see [38] for numerical data. Duplicate points correspond to $\chi = 2D$ and $\chi = 3D$.

one of the two subspaces. This allows us to probe explicitly for specific excited states and, among others, calculate the single-particle gap. Our ability to study physical systems in both parity sectors with tensor networks stands in contrast to infinite PEPS [18,21,31,39] or other renormalization-based algorithms [34,50,51] which often reach higher precisions for ground-state calculations.

The ground-state energies of the respective parity sectors can be determined with high accuracy. Using the boundary matrix product state method with a truncation dimension χ that scales linear in the bond dimension D allowed us to use large values of D to increase the precision. Extrapolations in squared inverse bond dimension D^{-2} allow for uncertainties in the order of 10^{-3} with relatively low computational costs. Nevertheless, a higher accuracy is required before physically meaningful interpretations of the single-particle gap on large lattices are possible. The single-particle gap is, of course, only one out of many different order parameters of the Hubbard model on the honeycomb lattice [52] and it might be worthwhile investigating additional observables like the staggered magnetization or correlator functions.

Our simulations work reliably and efficiently for large lattices and away from half-filling, but not necessarily with large onsite interactions. In the latter case, simple update becomes increasingly unstable. We also find that degenerate ground states can significantly decrease the stability of the simulations. While the only feasible solution to the first problem seems to be the usage of full update, the latter problem can usually be alleviated by explicitly breaking the degeneracy.

ACKNOWLEDGMENTS

We thank S. Kühn, R. Orús, and M. Carmen Bañuls for helpful discussions and remarks. We acknowledge the Computer Center at DESY Zeuthen for the compute time. This work was funded, in part, through financial support by the Helmholtz Einstein International Berlin Research School in Data Science (HEIBRiDS), by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) and by the NSFC through the funds provided to the Sino-German Collaborative Research Center CRC 110 “Symmetries and the Emergence of Structure in QCD” (DFG Project-ID No. 196253076-TRR 110, NSFC Grant No. 12070131001).

APPENDIX A: FERMIONIC SWAP GATES

We introduced the fermionic SWAP gates in Sec. II C as additional tensors in the tensor network. In our case we never have to contract the SWAP gate with a tensor explicitly. Instead, we can always find a way to write the tensor network such that the gate is contracted with two indices of a tensor $T_{i_1 i_2}$, where any additional indices of T are ignored at the moment because they are irrelevant for the consideration. With two contracted indices we can write

$$T_{i_1 i_2} X_{j_1 j_2}^{i_1 i_2} = T_{i_1 i_2} \delta_{i_1 j_1} \delta_{i_2 j_2} S(i_1, i_2) \quad (\text{A1})$$

$$= T_{j_1 j_2} S(j_1, j_2). \quad (\text{A2})$$

Note that the indices j_1 and j_2 in the last term (A2) are not contracted. This can be realized by an element wise multipli-

cation of T with the matrix S . Any additional indices of T simply yield a multiplicity of the same operation.

Figure 19 shows the different possible contractions of a SWAP gate. Wiggly lines summarize any number of indices that are not explicitly relevant for a given contraction but cannot be neglected nevertheless. Arrows between diagrams show the transformation under contraction.

The transformations in Fig. 19 are the different contractions of a SWAP gate as written in Eq. (A2). It becomes clear that a SWAP gate does precisely what one would expect by interchanging two lines (best to be seen in the central panel). The SWAP gate is self-inverse, so the transformation works the same way by a single contraction in both directions. The figure seems very redundant, but the swapping of different line types comes with different S matrices. Thus, the first panel interchanging two physical indices translates to

$$S^{pp} = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & 1 & -1 & -1 \end{pmatrix}, \quad (\text{A3})$$

the second physical-internal panel yields

$$S^{pi} = \begin{pmatrix} 1 & \dots & 1 & 1 & \dots & 1 \\ 1 & \dots & 1 & 1 & \dots & 1 \\ 1 & \dots & 1 & -1 & \dots & -1 \\ 1 & \dots & 1 & -1 & \dots & -1 \end{pmatrix}, \quad (\text{A4})$$

and the twice internal one reads as

$$S^{ii} = \begin{pmatrix} 1 & \dots & 1 & 1 & \dots & 1 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 1 & \dots & 1 & 1 & \dots & 1 \\ 1 & \dots & 1 & -1 & \dots & -1 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 1 & \dots & 1 & -1 & \dots & -1 \end{pmatrix}. \quad (\text{A5})$$

APPENDIX B: FULL UPDATE TRUNCATION FOR FERMIONS

The operator $\exp(-\tau H_i)$ is applied to a pair of nearest neighbors in the imaginary-time evolution. This increases the bond dimension and requires a truncation. We want to find new triads with a link of fixed bond dimension D approximating the new state. This is depicted in Fig. 20. The bond dimension splits into an even- and an odd-parity part $D = D_e + D_o$. We keep the size of the parity sectors fixed for this truncation procedure.

The truncation shall be optimal in the sense that the difference

$$f = \|\Psi - \Psi'\|^2 = \langle \Psi | \Psi \rangle - \langle \Psi | \Psi' \rangle - \langle \Psi' | \Psi \rangle + \langle \Psi' | \Psi' \rangle \quad (\text{B1})$$

between the old state and the new state is minimal. Scalar products mean summation of the corresponding physical indices between the bra and the ket as well as a contraction of the wiggly lines with the rest of the tensor network N_{red} (see Fig. 7).

The optimization involves two triads with two parity sectors each. Therefore, we have to update four tensors. As an




FIG. 19. Contraction of a single SWAP gate with a tensor. The left figure shows the two physical indices, the middle represents a physical and an internal index, and the right the two internal indices.

example, we present the update of the even-parity part of the left tensor here. The other tensors follow correspondingly. We first calculate the derivative of our loss function f with respect to the tensor to be updated:

$$\frac{\partial f}{\partial \text{tensor}} = \langle \frac{\partial \Psi'}{\partial \text{tensor}} | \Psi' \rangle - \langle \frac{\partial \Psi'}{\partial \text{tensor}} | \Psi \rangle \quad (\text{B2})$$

$$= \text{Diagram 1} + \text{Diagram 2} - \text{Diagram 3} \quad (\text{B3})$$




FIG. 20. Truncation after the application of an imaginary-time evolution operator. Left: original triads with operator applied. Right: new triads with given bond dimension D on the link between them.

To find the minimum of f , we set the derivative to zero. This leads to a linear equation for the triad of interest. Its solution is

$$\text{Diagram 4} = \text{Diagram 5} - \text{Diagram 6}, \quad (\text{B4})$$

where M^{-1} is the inverse of M :

$$\text{Diagram 7} := \text{Diagram 8}, \quad (\text{B5})$$

$$\text{Diagram 9} = \text{Diagram 10} \cdot e. \quad (\text{B6})$$

We get the optimal new tensor according to Eq. (B4) while keeping the other triads fixed. All four tensors are updated sequentially this way. We then calculate the truncation error f and repeat the procedure iteratively until f converged.

- [1] R. Orús, *Ann. Phys.* **349**, 117 (2014).
- [2] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, *Phys. Rev. Lett.* **96**, 220601 (2006).
- [3] D. Poilblanc, *Phys. Rev. B* **96**, 121118(R) (2017).
- [4] M. Mambrini, R. Orus, and D. Poilblanc, *Phys. Rev. B* **94**, 205124 (2016).
- [5] J.-Y. Chen, L. Vanderstraeten, S. Capponi, and D. Poilblanc, *Phys. Rev. B* **98**, 184409 (2018).
- [6] L. Tagliacozzo, A. Celi, and M. Lewenstein, *Phys. Rev. X* **4**, 041024 (2014).
- [7] E. Zohar, T. B. Wahl, M. Burrello, and J. I. Cirac, *Ann. Phys.* **374**, 84 (2016).
- [8] P. Corboz, R. Orús, B. Bauer, and G. Vidal, *Phys. Rev. B* **81**, 165104 (2010).
- [9] B. Vanhecke, F. Verstraete, and K. V. Acoleyen, *arXiv:2104.10564*.
- [10] M. Bañuls, K. Cichy, J. Cirac, and K. Jansen, *J. High Energy Phys.* **11** (2013) 158.
- [11] M. F. Parsons, A. Mazurenko, C. S. Chiu, G. Ji, D. Greif, and M. Greiner, *Science* **353**, 1253 (2016).
- [12] P. A. Lee, N. Nagaosa, and X.-G. Wen, *Rev. Mod. Phys.* **78**, 17 (2006).
- [13] H. Tasaki, *J. Phys.: Condens. Matter* **10**, 4353 (1998).

- [14] D. P. Arovas, E. Berg, S. Kivelson, and S. Raghu, [arXiv:2103.12097](#).
- [15] F. F. Assaad and I. F. Herbut, *Phys. Rev. X* **3**, 031010 (2013).
- [16] P. Buividovich, D. Smith, M. Ulybyshev, and L. von Smekal, *Phys. Rev. B* **99**, 205434 (2019).
- [17] J. Ostmeyer, E. Berkowitz, S. Krieg, T. A. Lähde, T. Luu, and C. Urbach, *Phys. Rev. B* **102**, 245105 (2020).
- [18] P. Corboz, *Phys. Rev. B* **93**, 045116 (2016).
- [19] A. K. Geim and K. S. Novoselov, *Nat. Mater.* **6**, 183 (2007).
- [20] A. H. Castro Neto, F. Guinea, N. M. R. Peres, K. S. Novoselov, and A. K. Geim, *Rev. Mod. Phys.* **81**, 109 (2009).
- [21] F. Verstraete and J. I. Cirac, [arXiv:cond-mat/0407066](#).
- [22] F. Verstraete and J. I. Cirac, *Phys. Rev. A* **70**, 060302(R) (2004).
- [23] M. B. Hastings, *J. Stat. Mech.: Theory Exp.* (2007) P08024.
- [24] J. Eisert, M. Cramer, and M. B. Plenio, *Rev. Mod. Phys.* **82**, 277 (2010).
- [25] A. T. Sornborger and E. D. Stewart, *Phys. Rev. A* **60**, 1956 (1999).
- [26] G. Vidal, *Phys. Rev. Lett.* **98**, 070201 (2007).
- [27] P. Corboz and G. Vidal, *Phys. Rev. B* **80**, 165129 (2009).
- [28] H. C. Jiang, Z. Y. Weng, and T. Xiang, *Phys. Rev. Lett.* **101**, 090603 (2008).
- [29] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, *New J. Phys.* **16**, 033014 (2014).
- [30] B. Bruognolo, J.-W. Li, J. von Delft, and A. Weichselbaum, *SciPost Phys. Lect. Notes* **25**, 1 (2021).
- [31] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **101**, 250602 (2008).
- [32] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, *Phys. Rev. B* **90**, 064425 (2014).
- [33] F. Verstraete, V. Murg, and J. I. Cirac, *Adv. Phys.* **57**, 143 (2008).
- [34] H. H. Zhao, Z. Y. Xie, Q. N. Chen, Z. C. Wei, J. W. Cai, and T. Xiang, *Phys. Rev. B* **81**, 174411 (2010).
- [35] MATLAB versions 9.8+ (R2020a-21a), The Mathworks, Inc., Natick, Massachusetts, <https://www.mathworks.com/products/matlab.html>.
- [36] R. N. C. Pfeifer, G. Evenbly, S. Singh, and G. Vidal, [arXiv:1402.0939](#).
- [37] V. Murg, F. Verstraete, and J. I. Cirac, *Phys. Rev. A* **75**, 033605 (2007).
- [38] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevB.104.155118> for calculated values.
- [39] P. Corboz, P. Czarnik, G. Kapteijns, and L. Tagliacozzo, *Phys. Rev. X* **8**, 031031 (2018).
- [40] M. Cristoforetti, F. Di Renzo, A. Mukherjee, and L. Scorzato, *Phys. Rev. D* **88**, 051501(R) (2013).
- [41] M. Cristoforetti, F. Di Renzo, G. Eruzzi, A. Mukherjee, C. Schmidt, L. Scorzato, and C. Torrero, *Phys. Rev. D* **89**, 114505 (2014).
- [42] M. Ulybyshev, C. Winterowd, and S. Zafeiropoulos, *Phys. Rev. D* **101**, 014508 (2020).
- [43] A. Alexandru, G. Basar, P. F. Bedaque, G. W. Ridgway, and N. C. Warrington, *J. High Energy Phys.* **05** (2016) 053.
- [44] A. Alexandru, G. Basar, and P. Bedaque, *Phys. Rev. D* **93**, 014504 (2016).
- [45] J.-L. Wynen, E. Berkowitz, S. Krieg, T. Luu, and J. Ostmeyer, *Phys. Rev. B* **103**, 125153 (2021).
- [46] M. M. Rams, P. Czarnik, and L. Cincio, *Phys. Rev. X* **8**, 041033 (2018).
- [47] K. Levenberg, *Q. Appl. Math.* **2**, 164 (1944).
- [48] D. W. Marquardt, *J. Soc. Ind. Appl. Math.* **11**, 431 (1963).
- [49] Nonlinear Least-Squares Fitting, GNU Scientific Library, www.gnu.org/software/gsl/doc/html/nls.html.
- [50] Z. Y. Xie, J. Chen, M. P. Qin, J. W. Zhu, L. P. Yang, and T. Xiang, *Phys. Rev. B* **86**, 045139 (2012).
- [51] D. Kadoh and K. Nakayama, [arXiv:1912.02414](#).
- [52] J. Ostmeyer, E. Berkowitz, S. Krieg, T. A. Lähde, T. Luu, and C. Urbach, [arXiv:2105.06936](#).