

Numerical stabilization of entanglement computation in auxiliary-field quantum Monte Carlo simulations of interacting many-fermion systems

Peter Broecker and Simon Trebst

Institute for Theoretical Physics, University of Cologne, 50937 Cologne, Germany

(Received 23 September 2016; published 19 December 2016)

In the absence of a fermion sign problem, auxiliary-field (or determinantal) quantum Monte Carlo (DQMC) approaches have long been the numerical method of choice for unbiased, large-scale simulations of interacting many-fermion systems. More recently, the conceptual scope of this approach has been expanded by introducing ingenious schemes to compute entanglement entropies within its framework. On a practical level, these approaches, however, suffer from a variety of numerical instabilities that have largely impeded their applicability. Here we report on a number of algorithmic advances to overcome many of these numerical instabilities and significantly improve the calculation of entanglement measures in the zero-temperature projective DQMC approach, ultimately allowing us to reach similar system sizes as for the computation of conventional observables. We demonstrate the applicability of this improved DQMC approach by providing an entanglement perspective on the quantum phase transition from a magnetically ordered Mott insulator to a band insulator in the bilayer square lattice Hubbard model at half filling.

DOI: [10.1103/PhysRevE.94.063306](https://doi.org/10.1103/PhysRevE.94.063306)

I. INTRODUCTION

In statistical physics, quantum Monte Carlo (QMC) simulations are a mainstay of the numerical characterization of quantum many-body systems [1]. Revered for their ability to provide unbiased, numerically exact results even in the strong coupling regime, they provide valuable insight and theoretical guidance where analytical approaches such as mean-field or perturbative calculations often fail. World-line approaches such as the stochastic series expansion (SSE) [2] (used in combination with nonlocal update schemes [3–5]) or the continuous-time formulation of the worm algorithm [6] have allowed the study of unfrustrated bosonic systems at unprecedented numerical precision. For fermionic many-body systems, the situation is somewhat more delicate with the infamous fermion sign problem [7,8] thwarting the polynomial efficiency of the approach for a range of problems, i.e., its power-law scaling with system size and inverse temperature. Nevertheless, one should not overlook that there are a number of fermionic quantum many-body systems that do not suffer from the fermion sign problem and at the same time exhibit interesting physics arising solely from the fermionic nature of its constituents. This includes, for instance, the formation of superconductivity in microscopic models of nearly antiferromagnetic metals [9,10], the emergence of topological order in Kane-Mele-Hubbard models [11–13], or fermionic quantum critical phenomena in Dirac matter [14–18] or Z_2 gauge field coupled fermion systems [19,20]. Large-scale lattice models of all these fermionic systems have been simulated using DQMC simulations in the past three years, typically with an emphasis on the statistical measurement of observables such as order parameters, correlation functions, or topological invariants. At the same time, conceptual advancements have paved the way to calculate entanglement measures such as Rényi entropies or entanglement spectra within the framework of DQMC simulations [21–24], in step with similar developments for their bosonic world-line counterparts [25,26]. However, the practical use of these entanglement computations has so far been impeded by a number of numerical instabilities arising

deep within the DQMC technique. The purpose of the paper at hand is to give a detailed account as to how to overcome these numerical insufficiencies and bring the computational effort of entanglement measures on a similar level with that of conventional observables. Our algorithmic improvements focus on the calculation of Rényi entropies using an adaptation of the replica scheme [27,28] to the DQMC approach [22]. The numerical stabilization of its computation includes a variety of algorithmic steps ranging from linear algebra aspects such as the inversion of seemingly singular matrices and technical tricks such as the inclusion of an artificial chemical potential, which is shown to control the condition numbers of the underlying linear algebra algorithms, to more DQMC specific aspects such as the optimal choice of Hubbard-Stratonovich transformation and respective auxiliary fields. Beyond the numerical stabilization we provide a detailed discussion of the convergence behavior of our enhanced algorithm. Finally, to illustrate the performance of the technical improvements laid out in this paper, we apply our DQMC+replica scheme framework to provide an entanglement perspective on the quantum phase transition between an antiferromagnetically ordered Mott insulator and a featureless band insulator in the bilayer Hubbard model on the square lattice.

The remainder of the paper is organized as follows. We start with a brief review of the general auxiliary-field quantum Monte Carlo approach in Sec. II where we put a particular emphasis on those aspects where numerical instabilities typically arise and how they are dealt with in the conventional algorithm [29–31] that is used to compute correlation functions at finite temperatures. We then turn to the replica scheme and discuss its conceptual adaptation to the DQMC framework in Sec. III. The main results of the paper are provided in Sec. IV where we discuss the numerical stabilization for the calculation of Rényi entropies, followed by a discussion of the optimal choice of Hubbard-Stratonovich transformation in Sec. IV C and the overall convergence behavior of the algorithm in Sec. IV D. The application of the thus-improved DQMC approach to the quantum phase transition in the bilayer

Hubbard model is presented in Sec. V. We close with a discussion in Sec. VI.

II. AUXILIARY-FIELD QUANTUM MONTE CARLO

We start our discussion with a brief recapitulation of the basics of the DQMC method [29]. We will primarily discuss its finite-temperature formulation and only touch on some of the aspects of the alternative zero-temperature projective scheme. This choice of focus is motivated by the fact that the algorithm for the computation of ground-state entanglement properties is conceptually much closer to the finite-temperature approach than the projective scheme. Readers looking for a more exhaustive introduction to DQMC approaches are referred to the review articles in Refs. [30,31].

A. Conceptual formalism

From a conceptual perspective, the goal of determinantal quantum Monte Carlo approaches is to sample the partition function

$$Z = \text{Tr} e^{-\beta \hat{H}} \quad (1)$$

for a given lattice Hamiltonian \hat{H} in the grand-canonical ensemble and to evaluate finite-temperature expectation values of correlation functions along the way. The first step in accomplishing this is to perform a Trotter-Suzuki decomposition, discretizing the inverse temperature or imaginary time into N_τ time slices of size $\Delta\tau = \beta/N_\tau$. The investigated Hamiltonians are typically of the form $\hat{H} = \hat{K} + \hat{V}$ with a noninteracting, quadratic part \hat{K} and an interacting, quartic part \hat{V} , respectively. By making use of the Baker-Campbell-Hausdorff formula, one can split the sum in the exponential into a product of exponentials and keep only the lowest-order term

$$Z = \text{Tr} \prod_{\tau=1}^{N_\tau} e^{-\Delta\tau \hat{K}} e^{-\Delta\tau \hat{V}} + O(\Delta\tau^2). \quad (2)$$

Ignoring the higher orders introduces an error of size $O(\Delta\tau^2)$ that can be considered negligible if $\Delta\tau$ is sufficiently small. To evaluate the interaction part, we apply a Hubbard-Stratonovich transformation [32,33] on each time slice and each term in the interaction operator reducing quartic to quadratic operators at the cost of introducing an auxiliary field that we denote by $s(\tau, j)$, with τ being the time slice and j the index indicating the decoupled term of the interaction operator. After this Hubbard-Stratonovich transformation, we end up with a system of free fermions coupled to the auxiliary field via an interaction $V(s, \tau)$ that depends on the auxiliary-field configuration and the time slice in imaginary time. The operators \hat{K} and $\hat{V}(s, \tau)$ are now simply one-particle operators and can be written out in their matrix form \mathbf{K} and $\mathbf{V}(s, \tau)$, respectively, so that the partition sum is given as

$$Z = \sum_{s(\tau, j)} \text{Tr} \prod_{\tau=1}^{N_\tau} e^{-\Delta\tau \mathbf{K}} e^{-\Delta\tau \mathbf{V}(s, \tau)} = \sum_{s(\tau, j)} \text{Tr} \prod_{\tau=1}^{N_\tau} \mathbf{B}(s, \tau), \quad (3)$$

where $\mathbf{B}(s, \tau)$ denotes the matrix product of the hopping matrix and the decoupled interaction matrix for a given time slice τ and a given auxiliary-field configuration s . Let us

introduce some notation that we will use throughout the following. First, in favor of a clearer notation, we will drop the auxiliary-field argument s , since we will focus on calculating objects for a fixed configuration. We define a partial product of slice matrices starting at time slice τ and ending at time slice τ' as

$$B(\tau', \tau) = \prod_{i=\tau}^{\tau'} \mathbf{B}(i), \quad (4)$$

The full product over all time slices is denoted as

$$B(\tau = 1) = \prod_{i=1}^{N_\tau} \mathbf{B}(i), \quad (5)$$

where the label τ in the latter indicates the starting point for the matrix product. If $\tau \neq 1$, then the corresponding matrix product should be viewed as cyclically permuted by $\tau - 1$ elements, a perspective that we will later use when accessing equal-time Green's functions at different imaginary time slices. Using this notation, we may write the partition sum in the following concise way

$$Z = \sum_{\{s(\tau, j)\}} \text{Tr} B(\tau = 1). \quad (6)$$

At this point, it should be noted that the partition sum still contains a trace over all fermion states and over all auxiliary-field configurations $s(\tau, j)$. We can, however, integrate out the free fermions, resulting in one determinant per auxiliary field and fermion configuration. It can further be shown [30,31] that the sum of determinants reduces to a single determinant quantifying the electron contribution to the partition sum for any given configuration of the auxiliary field. This allows us to reformulate the original partition sum (1) in terms of determinants as

$$Z = \sum_{\{s(\tau, j)\}} \det[1 + B(\tau)]. \quad (7)$$

A central object of quantum statistical physics is the equal-time Green's function $G(\tau)$, which is easily accessible in the DQMC algorithm with its matrix form related to the above determinants via

$$G(\tau) = [1 + B(\tau)]^{-1}. \quad (8)$$

The equal-time Green's function not only provides valuable access to correlation functions, it also plays a crucial role in devising an efficient updating procedure of the auxiliary-field configurations sampled at the heart of the Monte Carlo approach. For details we again refer to the literature [30,31].

B. Numerical implementation

Any practical implementation of the algorithm described above typically faces two problems, both related to the underlying linear algebra: (i) the evaluation of the matrix product (5) is found to be numerically unstable [34] and (ii) the resulting matrix $B(\tau)$ turns out to be ill conditioned, which hampers both the calculation of the Monte Carlo transition probabilities as well as the evaluation of the equal-time Green's function and any observables derived from it. The condition number of the matrix product increases with increasing hopping strength and

increasing inverse temperature, respectively, and may increase or decrease with increasing interaction strength depending on the type of Hubbard-Stratonovich transformation.

Stabilizing a matrix product of the form (5) is a well-understood procedure in numerical mathematics. The idea is to perform either a rank revealing QR or a singular value decomposition (SVD) after a certain number of multiplications m , typically chosen small enough such that this shorter product remains exact up to machine precision. By making use of one of the decompositions one recasts a given matrix M into a product $M = UDT$, where U is typically unitary, D is a diagonal matrix that contains the information about the inherent scales of the matrix (e.g., in the form of its singular value spectrum), and T is either unitary (SVD) or triangular (QR decomposition). Independent of the specific decomposition algorithm, we will refer to the values of the diagonal matrix D as the singular values. To stabilize the computation of the matrix product (5), we divide the imaginary time interval into N_m groups of m time slices representing a segment of length $\Delta = m \cdot \tau$ in imaginary time that encompasses m slice matrices $\mathbf{B}(i)$. After the multiplication and decomposition of the first m matrices into $U_1 D_1 T_1$, the next set of matrices is then multiplied only to U_1 , before D_1 is multiplied from the right and the resulting product is decomposed again, i.e.,

$$\left[\left(\prod_{i=m+1}^{2m} \mathbf{B}(i) \cdot U_1 \right) D_1 \right] T_1 = U_2 D_2 (T T_1) = U_2 D_2 T_2. \quad (9)$$

This procedure is repeated until all slice matrices are incorporated into the UDT decomposition. Note that we will use all the incomplete, intermediate results at a later point so they will be stored in memory. We refer to this set of matrices, $\{\{U_i\}, \{D_i\}, \{T_i\}\}$, as the stack.

The determinant that gives the weight of a specific configuration of the auxiliary field is given as

$$\det[\mathbb{1} + B(\tau)] = \det(\mathbb{1} + UDT), \quad (10)$$

and the equal-time Green's function as

$$G(\tau) = [\mathbb{1} + B(\tau)]^{-1} = (\mathbb{1} + UDT)^{-1}. \quad (11)$$

The addition of the identity matrix, although seemingly simple, requires careful attention to ensure an accurate calculation of the determinant and the Green's function. This problem is solved by separating the diagonal matrix D from the auxiliary matrices U and T before the resulting matrix is decomposed yet again:

$$(\mathbb{1} + UDT) = U(U^{-1}T^{-1} + D)T = (UU')D'(T'T). \quad (12)$$

The weight, i.e., the determinant, can then easily be read off as the product of the diagonal entries of D' . It should be noted that in practice one would calculate the logarithm of the determinant to avoid numerical overflow and rounding errors.

We already mentioned that the Green's function at a given imaginary time τ is also used in the updating procedure of the auxiliary field for the corresponding time slice, which means that we need to have access to the Green's function at each time slice. One option would be to recalculate the Green's function at every time slice from scratch, but that is numerically very

expensive. Instead, one may propagate the Green's function along imaginary time by virtue of

$$G(\tau + 1) = \mathbf{B}(\tau)G(\tau) \mathbf{B}^{-1}(\tau). \quad (13)$$

This equation suggests that we have to calculate the Green's function only once, e.g., at $\tau = 1$, and then continue propagating the Green's function along imaginary time updating the appropriate auxiliary-field degrees of freedom. However, the propagation suffers from the same problem as the buildup of the matrix product (5) itself and is stable only for a few steps. In practice, we assume the number of stable propagation steps N_{prop} to be equal to the number of slices m making up groups of \mathbf{B} matrices. In general, this number can be determined by considering the difference between the propagated Green's function and the recalculated one after m propagations have taken place. The elementwise relative deviation

$$\delta_{ij} = 2 \cdot \frac{|G_{ij}^{\text{recalculated}} - G_{ij}^{\text{propagated}}|}{|G_{ij}^{\text{recalculated}} + G_{ij}^{\text{propagated}}|} \quad (14)$$

between the two should be small enough, such that the update probabilities and measurements are not seriously affected. Therefore, one typically focuses on comparing only those matrix elements that are used in the updating and measurement procedure. For those imaginary time slices where the Green's function has to be recalculated from scratch, one can greatly facilitate its calculation by making use of the matrix products saved in the stack. We summarize the basic DQMC algorithm in the flow diagram provided in Fig. 1.

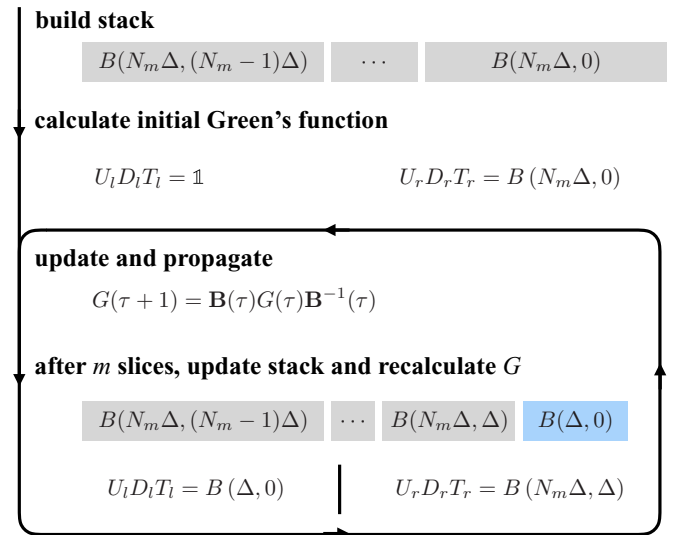


FIG. 1. Flow diagram of the basic DQMC algorithm. Shaded in gray is the initial buildup of the stack of decomposed matrix products. This is used to calculate the Green's function, which is subsequently updated and propagated in imaginary time before updating the stack and recalculating the Green's function from the stack to ensure stability. The slice product $B(\Delta, 0)$ shaded in blue has to be recalculated from scratch because it contains auxiliary-field values that were updated in the Monte Carlo procedure. Subsequent slice products will use the respective preceding U matrix as a starting point for the recalculation.

III. REPLICA SCHEME

To characterize and quantify the entanglement in a many-fermion system, we consider entanglement entropies [35,36] calculated for a bipartition of the system into two complementary subsystems A and B . Entanglement entropies have become a standard measure of entanglement for quantum many-body systems that can reveal valuable information about the principal nature of a quantum state [37]. Such a classification of quantum states is typically achieved by a careful scaling analysis of the entanglement entropies with system size, which reveals a dominant scaling with the length of the boundary of the bipartition [38] and several subleading contributions revealing, e.g., the formation of macroscopic, long-range entanglement in topologically ordered systems [39,40], the formation of more conventionally symmetry-broken ordered states via a detection of the corresponding Goldstone modes [41,42], or the existence of a Fermi surface via a logarithmic violation of the boundary law [43,44]. Technically, entanglement entropies are typically calculated in the form of Rényi entropies, which can be computed via the so-called replica scheme, a well-known procedure originally introduced in the context of quantum field theories [27]. In the following, we will quickly review how the replica scheme can be adapted to the framework of DQMC techniques.

A. Rényi entropies

Rényi entropies constitute a family of entanglement entropies calculated from the reduced density matrix $\rho_A = \text{Tr}_B \rho$ for a bipartition of the system into two complementary subsystems A and B

$$S_n = \frac{1}{1-n} \log \text{Tr}(\rho_A^n), \quad (15)$$

where the index n is typically an integer number. This family of Rényi entropies includes the well-known von Neumann entropy $S = -\text{Tr}(\rho_A \log \rho_A)$, which is recovered in the limit of $n \rightarrow 1$. For the purpose of our numerical Monte Carlo simulations we will, however, consider only Rényi entropies with $n > 1$, typically considering the case of the second Rényi entropy $n = 2$, for which one can formulate the replica scheme [27].

To quickly motivate and rederive this scheme let us start with the density matrix ρ of the quantum many-body system at hand. In general, we can think of this density matrix as being subject to a normalization via a factor $\text{Tr} \rho$, which is, of course, unaffected if one takes the partial trace over degrees of freedom of part B . Inserting this into Eq. (15) gives

$$S_n = \frac{1}{1-n} \log \frac{\text{Tr}(\rho_A^n)}{(\text{Tr} \rho)^n}, \quad (16)$$

which we rewrite as the ratio of two partition sums

$$S_n = \frac{1}{1-n} \log \frac{Z[A, n, \beta]}{Z^n}, \quad (17)$$

with $Z[A, n, \beta] = \text{Tr}(\rho_A^n)$ and $Z^n = (\text{Tr} \rho)^n$, respectively. In the following, let us focus on the second Rényi entropy, i.e.,




FIG. 2. World-line representation of the two partition sums of (18). On the left-hand side, the squared, regular partition function is depicted. World lines of particles are β periodic, no matter what subsystem they originate from. The right-hand side shows the modified partition sum that is β periodic in part B but 2β periodic in part A .

$n = 2$, for which the two equations above reduce to

$$S_2 = -\log \frac{\text{Tr}(\rho_A^2)}{(\text{Tr} \rho)^2} = -\log \frac{Z[A, 2, \beta]}{Z^2} \equiv -\log \frac{Z_1}{Z_0}. \quad (18)$$

Note that Z is the usual partition function, which we sample in conventional DQMC simulations to obtain numerical estimates of conventional observables such as correlation functions. The second partition function $Z[A, 2, \beta]$ entering the above equation is a modified partition function, where part B was traced out and the resulting reduced density matrix was squared ($n = 2$) before the final trace over A was taken. Associated to these partition functions are Green's functions G_0 and G_1 , respectively. Written out explicitly $Z[A, 2, \beta]$ takes the form

$$\begin{aligned} \text{Tr}_A(\rho_A^2) &= \sum_{A, A', B, B'} \langle AB' | \rho | A'B \rangle \langle A'B | \rho | AB \rangle \\ &\equiv Z[A, 2, \beta]. \end{aligned} \quad (19)$$

Note that $Z[A, 2, \beta]$ is β periodic in part B , which due to the squaring appears twice, once as B and once as B' , while part A is 2β periodic and appears only once. These peculiar boundary conditions in imaginary time can best be understood when visualized in a world-line picture as shown in Fig. 2.

B. Replica scheme in DQMC simulations

To adapt this replica scheme in DQMC simulations, we have previously shown [22] that one can fold out subsystems B and B' and consider a modified Hamiltonian that explicitly depends on imaginary time and takes the form

$$\tilde{H}(\tau) = \hat{H}_{AB} \Theta(\tau) \Theta(\beta - \tau) + \hat{H}_{AB'} \Theta(\tau - \beta) \Theta(2\beta - \tau). \quad (20)$$

Note that the simulation of this modified Hamiltonian also requires a modified simulation cell, where, according to Eq. (20), the degrees of freedom in subsystem A interact with those in B and B' not over the entire span of imaginary time but only in certain intervals as visualized in Fig. 3. We will refer to the subsystem B or B' that is currently in




FIG. 3. Modified simulation cell used in DQMC simulations. Part B appears twice as B and B' and both extend from 0 to 2β but its degrees of freedom only propagate in an interval of length β .

interaction with A as the active subsystem. In practice, the matrices associated with hopping and interaction processes are of size $(N_A + 2 \cdot N_B) \times (N_A + 2 \cdot N_B)$ where operators whose support lies exclusively in the inactive subsystem are represented by identity matrices and those mixing two subsystems are zero.

The replica scheme, as introduced above, has been formulated in the context of finite-temperature simulations with the inverse temperature β entering as a natural parameter in the replica scheme. We now want to adopt this approach to the zero-temperature projective DQMC algorithm that aims at directly sampling the ground-state wave function (in lieu of the partition sum). As will become clear in the following, this adaptation will result in an algorithm that is much closer to the finite-temperature scheme introduced above than the original ground-state algorithm [45]. At the heart of the ground-state algorithm is a projective scheme applied to a trial wave function $|\psi_T\rangle$

$$|\psi\rangle = \lim_{\theta \rightarrow \infty} e^{-\theta \hat{H}} |\psi_T\rangle, \quad (21)$$

which for sufficiently long projection times Θ yields the desired ground-state wave function $|\psi\rangle$. The only requirement imposed on the trial wave function is for it to have a nonzero overlap with the actual ground-state wave function, in which case the projection will eliminate all contributions from excited states and converge to the ground-state wave function in the limit of $\theta \rightarrow \infty$. Again, we can introduce a normalization constant—similar to the normalization of the density matrix in the discussion preceding Eq. (16) above, which can again be written as a trace over the ground-state density matrix

$$N = \langle \psi | \psi \rangle = \sum_{|A\rangle} \langle \psi | A \rangle \langle A | \psi \rangle = \text{Tr}(|\psi\rangle \langle \psi|). \quad (22)$$

We can readily insert this normalization constant (or trace) into the denominator of the definition of the Rényi entropy in Eq. (18). Inserting the projection (21) into the nominator of the definition of the Rényi entropy in (18), we find an expression for the $\text{Tr} \rho_A'^2$ that looks very similar to the finite-temperature

expression discussed above

$$\begin{aligned} \text{Tr} \rho_A'^2 &= \lim_{\theta \rightarrow \infty} \sum_{A, A', B, B'} \langle AB' | \exp(-\theta \hat{H}) |\psi_T\rangle \langle \psi_T| \\ &\quad \times \exp(-\theta \hat{H}) |A'B'\rangle \langle A'B| \exp(-\theta \hat{H}) |\psi_T\rangle \langle \psi_T| \\ &\quad \times \exp(-\theta \hat{H}) |AB\rangle, \end{aligned} \quad (23)$$

At first look, the only difference to Eq. (19) appears to be the occurrence of the density matrices $|\psi_T\rangle \langle \psi_T|$. However, note that the finite-temperature algorithm works in a grand-canonical ensemble, while the trace above is given in a canonical ensemble. This choice of ensembles is of crucial importance for the finite-temperature algorithm to work, because only then is the reduction to a sampling purely in the space of auxiliary-field configurations as shown in Eq. (7) possible. Nevertheless, it is possible to simply extend the canonical sums in Eq. (23) to grand-canonical ones because the density matrices act as projectors onto the correct particle sectors. Any combination of states that does not have the correct particle number will thus not contribute to the sum.

IV. NUMERICAL STABILIZATION

With the conceptual framework for the computation of Rényi entropies in DQMC simulations laid out in Secs. II and III above, we now turn to the technical aspects of its implementation. In particular, we discuss the numerical stabilization procedures enabling us to compute expectation values for these entanglement entropies for similar system sizes as for conventional observables (such as correlation functions) using the ordinary DQMC approach.

A. Invertibility of the matrix products

A fundamental problem in the adoption of the finite-temperature algorithm for ground-state properties is that the density matrix typically turns out to be singular, i.e., it is in general not invertible. As a consequence, the entire matrix product $B(\tau)$ and its decomposed form (11) are also no longer invertible, which ultimately prevents the computation of the equal-time Green's function and the associated Monte Carlo weights. However, one can overcome this algebraic obstacle by numerically constructing a matrix that is equal to the matrix product $B(\tau)$ up to machine precision, but whose matrix decomposition still yields invertible matrices as discussed in the following.

To see that the density matrix is in general a singular, noninvertible matrix let us consider a scenario where N_p particles populate a lattice of N sites. A trial wave function for such a scenario is typically constructed by diagonalizing the quadratic part of the Hamiltonian and then keeping only as many eigenvectors as the number of particles N_p in the system [46]. This test wave function is thus represented as a matrix ψ_T of dimension $N \times N_p$ with orthogonal columns, whereas the associated density matrix ρ of size $N \times N$ is constructed as $\rho = \psi_T \psi_T^\dagger$. From this construction it becomes apparent that the density matrix must be rank deficient (since it is the product of two nonsquare matrices) and as such singular.

To elude this problem, we modify the matrices U, D , and T of the singular value decomposition (12) such that they



(a) singular value decomposition (SVD)

(b) QR decomposition


FIG. 4. Inversion scheme as applicable to SVD (top) and QR (bottom) based Green’s function calculations. Shown here is the example of the projected trial density matrix $B(\tau, \tau/2)|\psi_T\rangle\langle\psi_T|B(\tau/2, 0)$. Matrices with orthogonal columns or rows that form an incomplete basis are extended to a full basis using a Gram-Schmidt process, denoted here in blue. The singular value vectors are extended by a vector of decreasing numbers clearly separated from the smallest physical singular value by the singular value gap Δ here pictured in green. For the case of the QR decomposition matrices without orthogonal columns or rows appear naturally, which are extended by a random matrix to match dimensions and ensure invertibility (for example by LU decomposition), pictured in orange.

become invertible but leave the original matrix unchanged (up to machine precision, which we denote by ϵ). In Fig. 4, we illustrate how to do this in practice: Any nonsquare matrix with orthogonal (unitary) columns (or rows) can be extended to a square, fully orthogonal (unitary) matrix by applying a Gram-Schmidt process. To ensure that the original matrix, i.e., the one that results from remultiplying $U D T$, remains unchanged, we extend the diagonal matrix by values that are at least $1/\epsilon$ smaller in magnitude than the smallest singular value found in the original D matrix. Numerically, this means that none of the additional columns or rows actually contribute, because they are weighted below numerical precision by the diagonal matrix. If we multiply the modified matrices $U \times D \times T$, we recover the original, singular square matrix B .

B. Stable calculation of the Green’s function

We now turn to the stable calculation of the equal-time Green’s function, which is calculated from the inverse of the matrix product $B(\tau)$, see Eq. (11), and therefore also sensitively depends on the matrix decomposition discussed in the previous section. Despite having circumvented the problem of singularity, the calculation of the Green’s functions remains difficult, because the matrices of the decomposition typically remain ill conditioned, i.e., they retain an extremely

broad singular value spectrum resulting in a high condition number [47]. The problem of calculating the Green’s function from ill-conditioned matrices in the finite-temperature algorithm has long been known and was solved by Hirsch and Fye [48] by using multiple consecutive matrix decompositions [instead of just one as given in Eq. (11)]. These multiple decompositions can be arranged in an enlarged matrix of size $(N_{\text{sites}} \cdot N_{\text{decompositions}}) \times (N_{\text{sites}} \cdot N_{\text{decompositions}})$, such that the determinant of this enlarged matrix remains equal to that of the original one, but the equal-time Green’s function can now be read off as a submatrix of the inverse of the enlarged matrix. The deeper reason that this approach allows to avoid the ill-conditioned matrix problem above is found in a considerably narrower singular value spectrum of the enlarged matrix. A maximum matrix size for the enlarged matrix is reached when each of the slice matrices is used as an input decomposition, i.e., $N_{\text{decompositions}} = N_\tau$. However, it is prohibitively expensive to invert such a large matrix [49]. Fortunately, for our entanglement computations at hand it typically suffices to choose just two or three consecutive decompositions for the calculation of G_0 (for a simulation cell with a complete cut) and up to five for G_1 (for a simulation cell with a partial cut).

One more technical caveat in the entanglement computation that warrants attention arises when inserting the density matrices $|\psi_T\rangle\langle\psi_T|$ in the calculation of the trace over the reduced density matrix in Eq. (23). Upon building a matrix product like $B = B(\theta, \theta/2)|\psi_i\rangle\langle\psi_i|B(\theta/2, 0)$ we start from the right and multiply slice matrices $\mathbf{B}(i)$, applying the stabilization procedure using successive matrix decompositions, until we built up the decomposition for the slice matrix group $B(\theta/2, 0)$. Up until this point in imaginary time, these matrices are square and invertible but now the insertion of the singular density matrix turns the entire product into a singular matrix. If we now carried on multiplying slice matrices on the left and decomposing the resulting matrices as we had done before, we would expect to find N_p nonzero singular values, corresponding to the particle number of the trial wave function and $N - N_p$ zero singular values. In practice, however, a decomposition algorithm like SVD or QR will typically find only the N_p nonzero singular values to high precision but the $N - N_p$ remaining singular values (strictly zero in theory) are found to be zero only relative to the actual nonzero singular values [47,49]. These inaccuracies will accumulate as we keep multiplying more slice groups and decomposing the resulting matrix products, which ultimately results in incorrect Green’s matrices. This problem can be overcome by a modification of the stack structure, whose implementation we discuss in the following for the case of calculating G_0 .

In this modified stack structure, we keep track of three such stacks: one that is built from the bra version of the wave function and includes decompositions obtained from matrices of the form $\langle\psi|B(\tau, \tau')$, a second one that is simply the ket version based on $B(\tau, \tau')|\psi\rangle$, and finally a third one that is used as temporary storage and includes decompositions of the full slice matrix groups $B(\tau, \tau')$. In combination, these three stacks allow us to calculate G_0 for a given imaginary time slice (at imaginary time $n \cdot \Delta$, see Sec. II B). An example configuration of the three stacks for imaginary time $\tau = n \cdot \Delta$ looks as follows.

stack 1

$$\langle \psi | B_m \quad \dots \quad \langle \psi | B_m B_{m-1} \dots B_i$$

stack 2

$$B_{2m} \dots B_{m+2} B_{m+1} | \psi \rangle \quad \dots \quad B_{m+1} | \psi \rangle$$

stack 3

$$B_{i-1} \dots B_1 \quad \dots \quad B_1$$

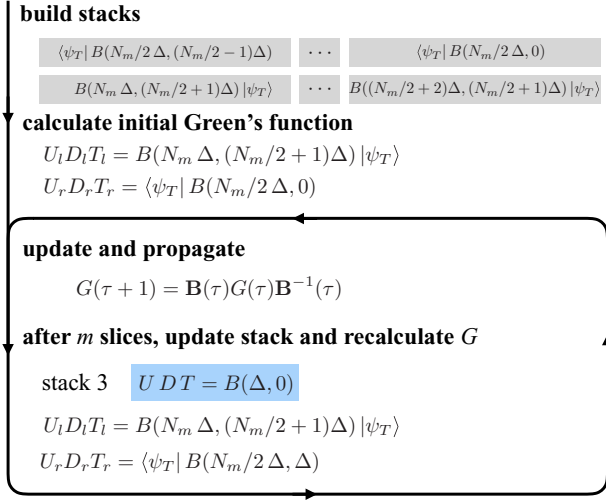


FIG. 5. Flow diagram of the modified DQMC algorithm for the simulation of the partition function Z_0 with complete cut in imaginary time. Stacks are initialized starting from the bra and ket version of the wave function and subsequently used to calculate an initial Green's function. This Green's function is then used to update the auxiliary-field time slice by time slice and recalculated after m time steps to retain numerical stability. Intermediate decompositions of slice matrix groups U, D, T without wave function are stored in a third stack, here highlighted by shading.

These three stacks can be used to set up a $3 \cdot N_{\text{sites}} \times 3 \cdot N_{\text{sites}}$ matrix and to subsequently calculate the Green's function and the corresponding Monte Carlo weight. Alternatively, if the condition numbers of the matrices allow it, it is possible to contract two or even all of them for a faster calculation. In that case, one has to think about the order of the matrix contractions. Naturally, the optimal choice is the one that keeps the condition numbers for the resulting matrices as small as possible, thus ensuring the highest stability of the following operations. Note that the number of contractions can vary for each recalculation and should be chosen according to the magnitude of the relative error between propagated and recalculated Green's function as defined in Eq. (14).

These numerical refinements are incorporated in the modified flow diagram of our DQMC algorithm in Fig. 5. As for the calculation of G_1 , we have to keep track of five individual stacks but the rest of the algorithm works exactly as just described for G_0 .

C. Choice of Hubbard-Stratonovich transformation

At the heart of the DQMC approach is the decoupling of quartic terms in the Hamiltonian using a Hubbard-Stratonovich

transformation. In general, there is a multitude of possible alternatives to perform this transformation. Depending on the physics of interest, common choices include a decoupling in the spin or charge channels, which leads to either a real or complex Hubbard-Stratonovich transformation that preserves or breaks $SU(2)$ symmetry, respectively. From an algorithmic point of view, the choice of Hubbard-Stratonovich transformation can greatly affect the numerical stability and convergence of the DQMC approach.

To illustrate the influence of the choice of Hubbard-Stratonovich transformation on the algorithmic performance, let us consider the example of a density-density interaction of the general form

$$\hat{U}_{\alpha\beta} = U_{\alpha\beta} n_{\alpha} n_{\beta}, \quad (24)$$

where the indices α and β may represent adjacent sites for a nearest-neighbor interaction or different spin species for a given site. In either case, we are looking to identify a quadratic operator A such that we can apply the Hubbard-Stratonovich transformation

$$e^{\frac{1}{2}A^2} = \sqrt{2\pi} \int_{-\infty}^{\infty} ds e^{-\frac{1}{2}s^2 - sA}, \quad (25)$$

where s is the auxiliary field introduced in the decoupling scheme. To do so, we will rewrite the interaction term (24) in an appropriate form. There are two possible choices

$$n_{\alpha} n_{\beta} = -\frac{1}{2}(n_{\alpha} - n_{\beta})^2 + \frac{1}{2}(n_{\alpha}^2 + n_{\beta}^2), \quad (26)$$

$$n_{\alpha} n_{\beta} = \frac{1}{2}(n_{\alpha} + n_{\beta})^2 - \frac{1}{2}(n_{\alpha}^2 + n_{\beta}^2), \quad (27)$$

with the quadratic operator corresponding to the first term, i.e., $A = (n_{\alpha} - n_{\beta})^2$ or $A = (n_{\alpha} + n_{\beta})^2$, respectively. For the scenario of a spin-spin on-site interaction the first choice of A in Eq. (26) corresponds to the squared magnetization and thus a decoupling in the spin channel. For the second choice (27) the operator A corresponds to the squared total charge and thus corresponds to a decoupling in the charge channel.

Note that in order to perform the Hubbard-Stratonovich transformation (25), the auxiliary field s does not necessarily have to be continuous, as probably suggested by the integral form of Eq. (25). For some cases, it suffices to work with a discrete field that is constrained to, e.g., take values ± 1 such as the case of the spinful Hubbard model, with most applications typically needing no more than a four-valued discrete field [50].

Applying the Hubbard-Stratonovich transformation (25) after bringing the operators into squared form (26, 27) gives the following expressions:

$$e^{-\Delta\tau U n_{\alpha} n_{\beta}} = \frac{1}{2} \sum_{s=\pm 1} \prod_{a=\alpha, \beta} e^{-(s\lambda + \frac{U\Delta\tau}{2})n_a} \quad (28)$$

and

$$e^{-\Delta\tau U n_{\alpha} n_{\beta}} = \frac{1}{2} \sum_{s=\pm 1} \prod_{a=\alpha, \beta} e^{-(s\lambda + \frac{U\Delta\tau}{2})(n_a - \frac{1}{2})}, \quad (29)$$

respectively. Here the constant λ is given by

$$\cosh(\lambda) = e^{\frac{1}{2}U\Delta\tau} \quad (30)$$

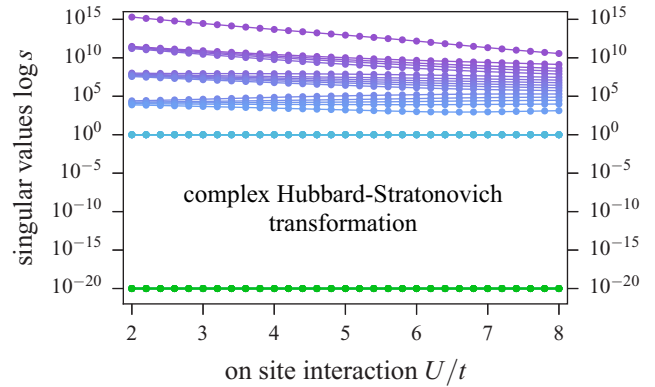



FIG. 6. The singular value spectrum of the matrix $B(\Theta/2,0)|\psi_T\rangle$, i.e., the projected wave function, on a logarithmic scale of base 10 for a real (left) and complex (right) Hubbard-Stratonovich transformation. The projection time is $\theta = 10$.

and

$$\cosh(\lambda) = e^{-\frac{1}{2}U\Delta\tau}, \tag{31}$$

respectively. Note that Eqs. (30), (31) imply that depending on the sign of the interaction U , the constant λ can either be real or complex valued.

At first look, one might think that the real-valued Hubbard-Stratonovich transformation leads to lower computational cost. However, in computations of the entanglement entropy it is the complex-valued Hubbard-Stratonovich transformation that should be used preferably. The reason for this preference is that the type of Hubbard-Stratonovich transformation also greatly affects the singular value spectrum of the matrix decompositions discussed in the previous Sec. IV B. To illustrate this point, we show in Fig. 6 a comparison of the singular value spectrum of the matrix $B(\Theta/2,0)|\psi_T\rangle$, i.e., the projected wave function, obtained for real and complex Hubbard-Stratonovich transformations. The data is calculated for a bilayer Hubbard model on a square lattice of size $4 \times 4 \times 2$ at half filling and equal hopping within and between the layers, i.e., $t = t'$, but varying on-site interaction U , see also Sec. V for a more detailed discussion of this model. In the complex case, the range of the singular values is seen to decrease as the interaction increases, hence reducing the condition number and stabilizing the algorithm. In the real case, the behavior is found to be exactly opposite with the condition number becoming worse as the interaction strength increases. As a consequence, we always use complex-valued Hubbard-Stratonovich transformation in the computation of Rényi entanglement entropies.

D. Convergence

One key distinction of our zero-temperature entanglement calculations in the projective DQMC approach outlined in this manuscript and the closely related finite-temperature algorithm is in the type of simulated ensemble. While the finite-temperature algorithm is sampling states from a grand-canonical ensemble, the projective scheme samples a canonical ensemble of fixed particle number (encoded in the trial wave function). This seemingly small modification of the ensemble is found to have some rather noticeable impact on the

convergence properties of the projective algorithm. To ensure convergence, we resort to a technical trick by introducing an artificial chemical potential in the Hamiltonian

$$H' = H + \mu_a \sum_i n_i. \tag{32}$$

Note that the inclusion of this artificial chemical potential μ_a is not altering the physics of the original problem, since we keep working in a canonical ensemble with a fixed particle number. To understand why it is nevertheless beneficial to include this term, let us first note that, on a technical level, such a term is represented by a diagonal matrix that has the ability to significantly shift the singular value spectrum. This is displayed in Fig. 7, which plots the smallest singular value of the unmodified matrix $B(\Theta/2,0)|\psi\rangle$. Clearly visible is the increase of the magnitude of the smallest singular value with increasing projection time Θ and increasing chemical potential μ_a . If we now consider the convergence of the




FIG. 7. Evolution of the smallest singular value with projection time Θ for various values for the artificial chemical potential μ_a shown on a logarithmic scale with base 10. The value of the smallest singular value increases with increasing artificial chemical potential μ_a .




FIG. 8. Convergence of the Rényi entanglement entropy S_2 in the projective DQMC algorithm versus projection time for varying values of the artificial chemical potential μ_a . The Rényi entropy is calculated for a half-filled Hubbard model on the square lattice for $U/t = 4$ with an equal-size bipartition of the lattice. For all nonzero artificial chemical potentials μ_a , the entanglement entropy S_2 converges to the same value with rate of convergence increasing from the smallest μ_a at the bottom to the largest μ_a at the top.

entanglement entropy upon inclusion of this artificial chemical potential term, we find that already a small additional potential dramatically improves the convergence of the entanglement entropy as a function of projection time as shown in Fig. 8. Further increasing the artificial chemical potential we find that the convergence eventually saturates, in the case at hand for about $\mu_a \approx 0.4$. In general, the optimal value of μ_a depends rather sensitively on the parameters of the Hamiltonian and needs to be chosen with great care.

Closer inspection of the convergence properties of our entanglement DQMC algorithm reveals that the added chemical potential (32) only affects the calculation of the Green's function G_1 associated with the partition sum Z_1 , i.e., the partition sum for the partially cut system, see Fig. 2. This observation already hints at the origin of the improved convergence. For the partition sum Z_1 all slice matrices are obtained from the unfolded Hamiltonian (20), where we fold out subsystem B as illustrated in Fig. 3. The slice matrices include blocks of identity and zero matrices to accommodate the currently inactive subsystem B or B' , i.e., the one which is deactivated by the Θ function in the Hamiltonian (20). As demonstrated in Fig. 9 this greatly affects the singular value spectrum by introducing N_B singular values of unit value (indicated by the magenta data set) representing the said subsystem. Comparing Figs. 7 and 9, we observe that once the difference between the smallest singular value of the active subsystem becomes separated from the unit singular values corresponding to the inactive subsystem, the simulation has a chance to converge. Thus, the requirement is that this gap Δ_s in the singular value spectrum has to be sufficiently large. In our studies we empirically find that $\Delta_s \approx 10^2$ turns out to be a good choice for all calculations.




FIG. 9. Magnitude of the singular values of the Green's function in the replicated system for a Hubbard model on a square lattice at half filling, an on site interaction of $U/t = 4$ and $\mu_a = 0.5$ on a logarithmic scale with base 10. The inactive subsystem in the middle and the invertibility noise on the bottom remain constant while the singular values of the active subsystems increase as projection time is increased.

V. APPLICATION TO THE BILAYER HUBBARD MODEL

We round off the technical discussion of this manuscript with one example illustrating the application of entanglement computations to identify quantum phases of interacting many-fermion systems. We resort to the bilayer Hubbard model—a paradigmatic model that allows us to study the transition between a Mott insulator and a band insulator. Whether these two insulating states are fundamentally distinct or can in general be adiabatically connected into one another [51] has been a question of debate [52–62]. This discussion mostly preceded the days of the topological insulator [63,64]—a second type of band insulator that can be clearly distinguished from the conventional trivial band insulator by certain topological invariants [65,66], while its fundamental distinction to correlated Mott insulators (possibly exhibiting topological order as well) is an open question of much current interest [67].

Here we want to apply an entanglement perspective on the elementary phase diagram of the bilayer Hubbard model

$$H = -t \sum_{(i,j),\sigma} c_{i,\sigma}^\dagger c_{j,\sigma} - t' \sum_{(i,j)',\sigma} c_{i,\sigma}^\dagger c_{j,\sigma} + U \sum_i n_{i,\uparrow} n_{i,\downarrow}, \quad (33)$$

which at its core describes the competition between conventional (free-fermion) band structures arising from hopping within and between the layer of a double-layer square lattice (parametrized by hopping amplitudes t and t' , respectively) and Mott physics arising from an on-site Coulomb repulsion U . A schematic phase diagram for this model in terms of the on-site interaction U and the interlayer hopping t' is given in Fig. 10. For sufficiently large interlayer hopping strength t' and any value of U the system is a featureless band insulator, while for sufficiently large U and small interlayer hopping t' the system forms a Mott insulator with antiferromagnetic




FIG. 10. Schematic phase diagram for the bilayer Hubbard model at half filling, adapted from dynamical mean-field theory (DMFT) [60], DQMC simulations [74], and functional renormalization group (FRG) calculations [75]. It comprises two insulating phases, a band insulator for sufficiently strong interlayer hopping $t' \gtrsim 4$, an antiferromagnetic insulator for weak interlayer hopping strength and large on-site Coulomb repulsion U . There possibly is an extended metallic phase in the regime of small Coulomb repulsion and interlayer hopping as indicated in the phase diagram, with DMFT and DQMC simulations in favor of it [60,74], while FRG calculations point to an absence of a metallic phase for any finite Coulomb repulsion [75]. The inset depicts the bilayer lattice. The intralayer hoppings are indicated by t , while t' denotes the interlayer hoppings.

order. The phase transition from Mott to band insulator has been studied using a variety of methods in the past [60,68–75]. In the absence of a Hubbard interaction U and small interlayer hopping we have a metal. Whether this metallic state survives for small Hubbard U (and small interlayer hopping t') or immediately gives way to Mott physics is still under debate, with dynamical mean-field theory (DMFT) [60] and finite-temperature DQMC simulations [70,74] pointing to an extended metallic phase, while a functional renormalization group (FRG) analysis in combination with a ground-state DQMC simulation [75] and a variational Monte Carlo study [73] find an immediate breakdown of the metallic phase upon inclusion of the Hubbard term.

Our DQMC approach is ideally suited to compute Rényi entanglement entropies over a wide range of parameters, in particular going deep into the regime of strong coupling $U > t, t'$. To do so, we consider a bipartition of the system into a strip of width $L/4 \times L \times 2$ and its complement. For this type of cut, which includes both the upper and lower layer, we expect that the entanglement signature of the band insulator is rather trivial—with the formation of singlet dimers on the rungs between the layers any cut between the rungs will effectively see no significant entanglement contribution, similar to the entanglement signature of an unentangled product state. Deep in the Mott regime we expect to observe the prevalent entanglement signature of any entangled quantum many-body state—a finite entanglement entropy that is subject to the famous boundary-law scaling [38]. For any metallic state we expect to observe an even stronger entanglement signature in the form of a logarithmic violation of the aforementioned boundary law [43,44].




FIG. 11. Color-coded Rényi entanglement entropy S_2 for the bilayer Hubbard model calculated for a bipartition of a $4 \times 4 \times 2$ system (with the system divided into a strip of extent $1 \times 4 \times 2$ and its complement). The entanglement entropy were calculated for integer values of U ranging from 1–10 and with t' in the range of $[0.5, 5.3]$ with a step size of 0.05. The different entanglement regimes clearly reveal the general characteristics of the underlying phase diagram. For dominant interlayer coupling $t' \gtrsim 4$ the vanishingly small entanglement points to the featureless band insulator with singlet formation on the interlayer rungs. For moderate interlayer coupling, the finite entanglement reveals the Mott insulator. The phase boundary separating the band insulator is indicated by an equipotential line of $S_2 = 0.8$. The relevant color values corresponding to these values of the entanglement entropy are also marked on the color bar by white lines.

As a first step we have scanned the absolute value of the Rényi entropy for $t' \in [0.5, 5.3]$ as illustrated in Fig. 11 for a system of $4 \times 4 \times 2$ sites. Clearly, several distinct regimes of almost constant amount of entanglement can be readily distinguished thereby revealing some of the core features of the phase diagram. For instance, one can clearly identify a regime of almost vanishing entanglement for large interlayer coupling t' —this is the band insulator as argued above. The finite entanglement clearly present in the lower part of the




FIG. 12. Rényi entanglement entropy S_2 for the bilayer Hubbard model along a vertical cut through the phase diagram of Fig. 11 at $U/t = 4$. Data for different system sizes is shown.




FIG. 13. Finite-size scaling of the entanglement entropy S_2 renormalized by the linear system size L for the bilayer Hubbard model at different values of the on-site interaction U/t and fixed interlayer coupling $t'/t = 1$. Note the logarithmic scale with base 10 chosen for the abscissa.

phase diagram points to at least one separate regime. The Rényi entropy along a vertical cut through the data of Fig. 11 is shown in Fig. 12 for $U/t = 4$ and $L = 4$ and $L = 8$, with the intermediate steplike behavior for the smaller system most likely a finite-size effect.

To further substantiate the nature of the phase(s) in the lower half of the phase diagram we consider the scaling of the Rényi entropy with system size. In Fig. 13 we show results for the Rényi entropy for different values of the on-site Coulomb repulsion $U/t = 2, 4, 8$, and 16 and fixed interlayer hopping $t'/t = 1$. With the entropy renormalized by the linear system size L , the boundary law scaling expected for the Mott insulator corresponds to a flat line. This is precisely what we find for moderate to large Hubbard interaction $U/t = 4, 8, 16$. In contrast, for small Hubbard interaction $U/t = 2$ the clearly noticeable slope of the data points suggests an additional logarithmic contribution to the entanglement entropy. Precisely such an $S \propto L \log L$ scaling is expected for a two-dimensional metal with a Fermi surface [43,44]. These results are consistent with finite-temperature data given for smaller lattices in Ref. [74]. Note, however, that we are looking at rather small system sizes when inferring this logarithmic contribution and it would be highly desirable to go to substantially larger system sizes to exclude possible finite-size effects, see Ref. [75] for a more in-depth discussion and analysis from the point of view of correlation functions. Unfortunately, such a substantial increase in system sizes does

not seem feasible with the current algorithms and state-of-the-art computational resources. Nevertheless, with the data at hand the entanglement point of view is certainly adding support to a scenario, in which the metallic state survives the inclusion of a moderately small Hubbard interaction.

VI. SUMMARY

To summarize, we have laid out a number of algorithmic advances to overcome the numerical instabilities of computing entanglement entropies for interacting many-fermion systems within the DQMC framework. The basic algorithm [22] is an adaptation of the replica scheme to the DQMC approach, which on a technical level converts the finite-temperature algorithm (in the grand-canonical ensemble) to a ground-state projective algorithm (in the canonical ensemble). Our numerical improvements of this approach, reported on here, include (i) a numerical stabilization of the matrix inversion of seemingly singular matrix products via a refined matrix decomposition, (ii) a numerical stabilization of the calculation of the Green's function using this refined matrix decomposition, (iii) an optimal choice of Hubbard-Stratonovich transformation, which typically turns out to be a complex-valued one, and (iv) the inclusion of an artificial chemical potential to significantly improve the convergence of the algorithm. With these algorithmic enhancements implemented, we find that one can reach similar system sizes as for the calculation of more conventional observables in the original projective DQMC scheme and showcased the algorithm by applying it to the bilayer Hubbard model.

From a physics point of view, probably the most interesting application of the computational scheme at hand is to showcase that one can positively identify the formation of macroscopic entanglement in an interacting many-fermion system exhibiting a phase with intrinsic topological order such as the recently proposed models of Dirac fermions coupled to a Z_2 gauge field [19,20]. This is ongoing work.

ACKNOWLEDGMENTS

We thank F. Assaad for insightful discussions on the technical aspects of the DQMC approach and S. Wessel for an exchange on the conflicting numerical results for the phase diagram of the bilayer Hubbard model. P.B. acknowledges partial support from the Deutsche Telekom Stiftung and the Bonn-Cologne Graduate School of Physics and Astronomy (BCGS). We thank the DFG for partial support within the CRC network TR 183 (Project B01). The numerical simulations were performed on the CHEOPS cluster at RRZK Cologne and the JURECA cluster at the Forschungszentrum Juelich.

- [1] J. Gubernatis, N. Kawashima, and P. Werner, *Quantum Monte Carlo methods: Algorithms for lattice models* (Cambridge University Press, Cambridge, 2016).
- [2] A. W. Sandvik, Stochastic series expansion method with operator-loop update, *Phys. Rev. B* **59**, 14157(R) (1999).

- [3] O. F. Syljuåsen and A. W. Sandvik, Quantum Monte Carlo with directed loops, *Phys. Rev. E* **66**, 046701 (2002).
- [4] F. Alet, S. Wessel, and M. Troyer, Generalized directed loop method for quantum Monte Carlo simulations, *Phys. Rev. E* **71**, 036706 (2005).

- [5] M. Troyer, F. Alet, S. Trebst, and S. Wessel, Non-local updates for quantum Monte Carlo simulations, *AIP Conf. Proc.* **690**, 156 (2003).
- [6] N. V. Prokof'ev, B. V. Svistunov, and I. S. Tupitsyn, Exact, complete, and universal continuous-time worldline Monte Carlo approach to the statistics of discrete quantum systems, *J. Exp. Theor. Phys.* **87**, 310 (1998).
- [7] J. E. Hirsch, R. L. Sugar, D. J. Scalapino, and R. Blankenbecler, Monte Carlo simulations of one-dimensional fermion systems, *Phys. Rev. B* **26**, 5033 (1982).
- [8] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, Sign problem in the numerical simulation of many-electron systems, *Phys. Rev. B* **41**, 9301 (1990).
- [9] E. Berg, M. A. Metlitski, and S. Sachdev, Sign-problem-free quantum Monte Carlo of the onset of antiferromagnetism in metals, *Science* **338**, 1606 (2012).
- [10] Y. Schattner, M. H. Gerlach, S. Trebst, and E. Berg, Competing Orders in a Nearly Antiferromagnetic Metal, *Phys. Rev. Lett.* **117**, 097002 (2016).
- [11] M. Bercx, M. Hohenadler, and F. F. Assaad, Kane-Mele-Hubbard model on the π -flux honeycomb lattice, *Phys. Rev. B* **90**, 075140 (2014).
- [12] M. Hohenadler, F. Parisen Toldin, I. F. Herbut, and F. F. Assaad, Phase diagram of the Kane-Mele-Coulomb model, *Phys. Rev. B* **90**, 085146 (2014).
- [13] F. Parisen Toldin, M. Hohenadler, F. F. Assaad, and I. F. Herbut, Fermionic quantum criticality in honeycomb and π -flux Hubbard models: Finite-size scaling of renormalization-group-invariant observables from quantum Monte Carlo, *Phys. Rev. B* **91**, 165108 (2015).
- [14] L. Wang, P. Corboz, and M. Troyer, Fermionic quantum critical point of spinless fermions on a honeycomb lattice, *New J. Phys.* **16**, 103008 (2014).
- [15] Z.-X. Li, Y.-F. Jiang, and H. Yao, Fermion-sign-free Majorana-quantum-Monte-Carlo studies of quantum critical phenomena of Dirac fermions in two dimensions, *New J. Phys.* **17**, 085003 (2015).
- [16] J. Motruk, A. G. Grushin, F. de Juan, and F. Pollmann, Interaction-driven phases in the half-filled honeycomb lattice: An infinite density matrix renormalization group study, *Phys. Rev. B* **92**, 085147 (2015).
- [17] S. Capponi and A. M. Läuchli, Phase diagram of interacting spinless fermions on the honeycomb lattice: A comprehensive exact diagonalization study, *Phys. Rev. B* **92**, 085146 (2015).
- [18] P. Broecker and S. Trebst, Entanglement and the fermion sign problem in auxiliary-field quantum Monte Carlo simulations, *Phys. Rev. B* **94**, 075144 (2016).
- [19] F. Assaad and T. Grover, A simple fermionic model of deconfined phases and phase transitions, [arXiv:1607.03912](https://arxiv.org/abs/1607.03912) (unpublished).
- [20] S. Gazit, M. Randeria, and A. Vishwanath, Charged fermions coupled to \mathbb{Z}_2 gauge fields: Superfluidity, confinement and emergent Dirac fermions, [arXiv:1607.03892](https://arxiv.org/abs/1607.03892) (unpublished).
- [21] T. Grover, Entanglement of Interacting Fermions in Quantum Monte Carlo Calculations, *Phys. Rev. Lett.* **111**, 130402 (2013).
- [22] P. Broecker and S. Trebst, Rényi entropies of interacting fermions from determinantal quantum Monte Carlo simulations, *J. Stat. Mech.: Theor. Expt.* (2014) P08015.
- [23] F. F. Assaad, T. C. Lang, and F. Parisen Toldin, Entanglement spectra of interacting fermions in quantum monte carlo simulations, *Phys. Rev. B* **89**, 125121 (2014).
- [24] F. F. Assaad, Stable quantum Monte Carlo simulations for entanglement spectra of interacting fermions, *Phys. Rev. B* **91**, 125146 (2015).
- [25] M. B. Hastings, I. González, A. B. Kallin, and R. G. Melko, Measuring Renyi Entanglement Entropy in Quantum Monte Carlo Simulations, *Phys. Rev. Lett.* **104**, 157201 (2010).
- [26] C. M. Herdman, S. Inglis, P.-N. Roy, R. G. Melko, and A. Del Maestro, Path-integral Monte Carlo method for Rényi entanglement entropies, *Phys. Rev. E* **90**, 013308 (2014).
- [27] C. Holzhey, F. Larsen, and F. Wilczek, Geometric and renormalized entropy in conformal field theory, *Nucl. Phys. B* **424**, 443 (1994).
- [28] P. Calabrese and J. Cardy, Entanglement entropy and quantum field theory, *J. Stat. Mech.: Theor. Expt.* (2004) P06002.
- [29] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, Monte Carlo calculations of coupled boson-fermion systems. I, *Phys. Rev. D* **24**, 2278 (1981).
- [30] R. R. dos Santos, Introduction to quantum Monte Carlo simulations for fermionic systems, *Braz. J. Phys.* **33**, 36 (2003).
- [31] F. F. Assaad and H. G. Evertz, in *Computational Many-Particle Physics*, Lecture Notes in Physics No. 739, edited by H. Fehske, R. Schneider, and A. Weiße (Springer, Berlin, 2008), pp. 277–356.
- [32] R. L. Stratonovich, On a method of calculating quantum distribution functions, *Soviet Physics Doklady* **2**, 416 (1957).
- [33] J. Hubbard, Calculation of Partition Functions, *Phys. Rev. Lett.* **3**, 77 (1959).
- [34] Z. Bai, C. Lee, R.-C. Li, and S. Xu, Stable solutions of linear systems involving long chain of matrix multiplications, *Linear Algebra Appl.* **435**, 659 (2011).
- [35] J. D. Bekenstein, Black holes and entropy, *Phys. Rev. D* **7**, 2333 (1973).
- [36] S. W. Hawking, Particle creation by black holes, *Commun. Math. Phys.* **43**, 199 (1975).
- [37] N. Laflorencie, Quantum entanglement in condensed matter systems, *Phys. Rep.* **646**, 1 (2016).
- [38] J. Eisert, M. Cramer, and M. B. Plenio, *Colloquium: Area laws for the entanglement entropy*, *Rev. Mod. Phys.* **82**, 277 (2010).
- [39] M. Levin and X.-G. Wen, Detecting Topological Order in a Ground State Wave Function, *Phys. Rev. Lett.* **96**, 110405 (2006).
- [40] A. Kitaev and J. Preskill, Topological Entanglement Entropy, *Phys. Rev. Lett.* **96**, 110404 (2006).
- [41] M. A. Metlitski and T. Grover, Entanglement entropy of systems with spontaneously broken continuous symmetry, [arXiv:1112.5166](https://arxiv.org/abs/1112.5166) (unpublished).
- [42] B. Kulchytskyy, C. M. Herdman, S. Inglis, and R. G. Melko, Detecting Goldstone modes with entanglement entropy, *Phys. Rev. B* **92**, 115146 (2015).
- [43] M. M. Wolf, Violation of the Entropic Area Law for Fermions, *Phys. Rev. Lett.* **96**, 010404 (2006).
- [44] D. Gioev and I. Klich, Entanglement Entropy of Fermions in Any Dimension and the Widom Conjecture, *Phys. Rev. Lett.* **96**, 100503 (2006).
- [45] S. Sorella, S. Baroni, R. Car, and M. Parrinello, A novel technique for the simulation of interacting fermion systems, *Europhys. Lett.* **8**, 663 (1989).

- [46] Note that special care has to be taken when the eigenvalues are degenerate [50]. In that case, one can for example randomize the hopping amplitudes a little bit to lift the degeneracy .
- [47] L. N. Trefethen and D. Bau, *Numerical Linear Algebra* (SIAM, Philadelphia, 1997).
- [48] R. M. Fye and J. E. Hirsch, Monte Carlo study of the symmetric Anderson-impurity model, *Phys. Rev. B* **38**, 433 (1988).
- [49] J. W. Demmel, *Applied Numerical Linear Algebra* (SIAM, Philadelphia, 1997).
- [50] F. F. Assaad, M. Imada, and D. J. Scalapino, Charge and spin structures of a $d_{x^2-y^2}$ superconductor in the proximity of an antiferromagnetic Mott insulator, *Phys. Rev. B* **56**, 15001 (1997).
- [51] F. Anfuso and A. Rosch, String order and adiabatic continuity of Haldane chains and band insulators, *Phys. Rev. B* **75**, 144420 (2007).
- [52] F. H. L. Essler and A. M. Tsvelik, Weakly coupled one-dimensional Mott insulators, *Phys. Rev. B* **65**, 115117 (2002).
- [53] F. H. L. Essler and A. M. Tsvelik, Theory of hybrid state in a metal with a small Fermi surface and strong collective excitations, *Phys. Rev. B* **71**, 195116 (2005).
- [54] A. Fuhrmann, D. Heilmann, and H. Monien, From Mott insulator to band insulator: A dynamical mean-field theory study, *Phys. Rev. B* **73**, 245118 (2006).
- [55] R. M. Konik, T. M. Rice, and A. M. Tsvelik, Doped Spin Liquid: Luttinger Sum Rule and Low Temperature Order, *Phys. Rev. Lett.* **96**, 086407 (2006).
- [56] C. Berthod, T. Giamarchi, S. Biermann, and A. Georges, Breakup of the Fermi Surface Near the Mott Transition in Low-Dimensional Systems, *Phys. Rev. Lett.* **97**, 136401 (2006).
- [57] A. Garg, H. R. Krishnamurthy, and M. Randeria, Can Correlations Drive a Band Insulator Metallic?, *Phys. Rev. Lett.* **97**, 046403 (2006).
- [58] T. D. Stanescu, P. Phillips, and T.-P. Choy, Theory of the Luttinger surface in doped Mott insulators, *Phys. Rev. B* **75**, 104503 (2007).
- [59] A. Rosch, Breakdown of Luttinger's theorem in two-orbital Mott insulators, *Eur. Phys. J. B* **59**, 495 (2007).
- [60] S. S. Kancharla and S. Okamoto, Band insulator to Mott insulator transition in a bilayer Hubbard model, *Phys. Rev. B* **75**, 193103 (2007).
- [61] N. Paris, K. Bouadim, F. Hebert, G. G. Batrouni, and R. T. Scalettar, Quantum Monte Carlo Study of an Interaction-Driven Band-Insulator to Metal Transition, *Phys. Rev. Lett.* **98**, 046403 (2007).
- [62] S. S. Kancharla and E. Dagotto, Correlated Insulated Phase Suggests Bond Order between Band and Mott Insulators in Two Dimensions, *Phys. Rev. Lett.* **98**, 016402 (2007).
- [63] M. Z. Hasan and C. L. Kane, *Colloquium*: Topological insulators, *Rev. Mod. Phys.* **82**, 3045 (2010).
- [64] X.-L. Qi and S.-C. Zhang, Topological insulators and superconductors, *Rev. Mod. Phys.* **83**, 1057 (2011).
- [65] C. L. Kane and E. J. Mele, Quantum Spin Hall Effect in Graphene, *Phys. Rev. Lett.* **95**, 226801 (2005).
- [66] C. L. Kane and E. J. Mele, Z_2 Topological Order and the Quantum Spin Hall Effect, *Phys. Rev. Lett.* **95**, 146802 (2005).
- [67] C. Wang, A. C. Potter, and T. Senthil, Classification of interacting electronic topological insulators in three dimensions, *Science* **343**, 629 (2014).
- [68] Z. Weihong, Various series expansions for the bilayer S=Heisenberg antiferromagnet, *Phys. Rev. B* **55**, 12267 (1997).
- [69] L. Wang, K. S. D. Beach, and A. W. Sandvik, High-precision finite-size scaling analysis of the quantum-critical point of S = 1/2 Heisenberg antiferromagnetic bilayers, *Phys. Rev. B* **73**, 014431 (2006).
- [70] K. Bouadim, G. G. Batrouni, F. Hébert, and R. T. Scalettar, Magnetic and transport properties of a coupled Hubbard bilayer with electron and hole doping, *Phys. Rev. B* **77**, 144527 (2008).
- [71] C. J. Hamer, J. Oitmaa, and Z. Weihong, Restoration of symmetry in the spectrum of the bilayer Heisenberg antiferromagnet, *Phys. Rev. B* **85**, 014432 (2012).
- [72] J. Helmes and S. Wessel, Entanglement entropy scaling in the bilayer Heisenberg spin system, *Phys. Rev. B* **89**, 245120 (2014).
- [73] R. Rüger, L. F. Tocchio, R. Valenti, and C. Gros, The phase diagram of the square lattice bilayer Hubbard model: A variational Monte Carlo study, *New J. Phys.* **16**, 033010 (2014).
- [74] C.-C. Chang, R. R. P. Singh, and R. T. Scalettar, Entanglement properties of the antiferromagnetic-singlet transition in the Hubbard model on bilayer square lattices, *Phys. Rev. B* **90**, 155113 (2014).
- [75] M. Golor, T. Reckling, L. Classen, M. M. Scherer, and S. Wessel, Ground-state phase diagram of the half-filled bilayer Hubbard model, *Phys. Rev. B* **90**, 195131 (2014).