Direct sampling of projected entangled-pair states

Tom Vieijra,^{*} Jutho Haegeman, Frank Verstraete, and Laurens Vanderstraeten

Department of Physics and Astronomy, Ghent University, B-9000 Ghent, Belgium

(Received 23 September 2021; revised 29 November 2021; accepted 30 November 2021; published 21 December 2021)

Variational Monte Carlo studies employing projected entangled-pair states (PEPS) have recently shown that they can provide answers to long-standing questions such as the nature of the phases in the two-dimensional J_1 - J_2 model. The sampling in these Monte Carlo algorithms is typically performed with Markov chain Monte Carlo algorithms employing local update rules, which often suffer from long autocorrelation times and interdependent samples. We propose a sampling algorithm that generates independent samples from a PEPS, bypassing all problems related to finite autocorrelation times. This algorithm is a generalization of an existing direct sampling algorithm for unitary tensor networks. We introduce an auxiliary probability distribution from which independent samples can be drawn, and combine it with importance sampling in order to evaluate expectation values accurately. We benchmark our algorithm on the classical Ising model and on variational optimization of two-dimensional quantum spin models.

DOI: [10.1103/PhysRevB.104.235141](https://doi.org/10.1103/PhysRevB.104.235141)

I. INTRODUCTION

Strong correlations in quantum many-body systems often lead to interesting phenomena, but are notoriously hard to understand with traditional many-body techniques such as perturbation theory or mean-field approaches. On the computational side, exact diagonalization is necessarily limited to small clusters and quantum Monte Carlo techniques often suffer from the sign problem. In many cases, the use of variational wave functions enables us to get insight into the nature of strongly correlated ground states. In order to decide on the ground-state properties for a given model Hamiltonian, however, we would like to consider variational classes with a (large) number of parameters that are variationally optimized. In the variational Monte Carlo (VMC) approach this optimization is performed stochastically as the variational energy is evaluated by a sampling procedure. This VMC approach can be successfully applied for optimizing Jastrow wave functions, for describing spin-liquid ground states [1], and optimizing variational states inspired by neural networks [2].

Tensor networks represent a particularly versatile class of variational states for describing low-energy states of quantum lattice systems. For one-dimensional systems, matrix product states (MPS) are extremely successful, because local expectation values and correlation functions can be computed exactly with low computational cost, and they can be optimized efficiently by the density-matrix renormalization group (DMRG) algorithm or direct variational optimization [3]. In two dimensions, projected entangled-pair states (PEPS) [4] are less efficient: evaluating expectation values comes at a significantly larger computational cost, and finding efficient schemes for the variational optimization is still ongoing [4–13]. Nonetheless, PEPS are, by now, routinely used as a

computational tool for simulating ground-state properties of strongly correlated 2-D quantum systems [14–17], and the framework is being extended to the simulation of thermal states [18–20], excited states [21–23], and time evolution [24–26].

In that respect, it is natural to consider the use of sampling techniques for evaluating and optimizing PEPS wave functions. Although there have been a few works that combine tensor networks with sampling techniques [27–32], only recently stochastic optimization has been put forward as an efficient optimization algorithm for PEPS [11,33]; this method showed very strong results for the frustrated J_1 - J_2 Heisenberg model on the square lattice [34]. In these works, the Markov chain Monte Carlo (MCMC) algorithm was applied to sample PEPS wave functions efficiently—as is standard in the VMC approach for simulating quantum systems.

Recently, however, interest has mounted in sampling procedures that do not involve a Markov chain, but instead are able to generate a sample directly from the probability distribution. A famous example of these so-called direct sampling approaches is the Box-Muller algorithm to sample variables distributed by Gaussians [35]. This approach has the great advantage that all samples can be generated in parallel, which is especially useful in light of recent advances in computer hardware such as massively parallel GPUs. Furthermore, the samples are independent from each other, which is not the case in MCMC methods. Related to this interdependency in MCMC are aspects that reduce its efficiency such as long autocorrelation times and poor mixing. By employing direct sampling schemes, these shortcomings are eliminated. Examples of recent progress in direct sampling are flow models [36,37] and autoregressive neural networks [38,39] in the context of generative machine learning. These have recently been used in a physical context, such as for sampling classical many-body systems [40], in the context of lattice field theories [41] and optimization of variational states for many-body systems [42–44].

^{*}tom.vieijra@ugent.be

In the case of MPS, the use of canonical forms can be used to construct a direct sampling algorithm, where the physical degrees of freedom are sampled according to the square of the wave function $|\Psi|^2$ [45]. The canonical form allows to easily construct low-dimensional conditional distributions, that can be sampled sequentially to generate a configuration from a high-dimensional probability distribution. Recently, the conditional distributions for classical partition functions modeled as two-dimensional tensor networks were approximated by a tensor renormalization algorithm and the resulting samples were subsequently used as trial states for a Metropolis algorithm [46]. Although canonical forms are not generic for PEPS representations and direct sampling according to $p_\Psi = |\Psi|^2$ is not possible, in this paper we propose an approximate scheme for generating samples efficiently. We construct an auxiliary distribution p_c , from which we can sample directly, and that is closely related to p_Ψ . The mismatch between p_c and p_Ψ can be systematically reduced, and is accounted for in the estimation of expectation values by employing importance sampling. We show that this scheme performs well, and we use this direct sampling approach to efficiently find variational ground states of two-dimensional interacting spin systems.

The paper is organized as follows. In Sec. II, we review the PEPS construction and how sampling can be used to estimate observables. In Sec. III, we describe the previously used MCMC algorithm, and we introduce our direct sampling algorithm. In Sec. IV, we introduce the details of the variational optimization algorithm we use to optimize PEPS wave functions. In Sec. V, we benchmark our sampling algorithm on the classical Ising model (which can be written as a simple PEPS) and on the variational optimization of the transverse-field Ising model and the antiferromagnetic Heisenberg model.

II. PROJECTED ENTANGLED-PAIR STATES

We consider quantum spin systems on a two-dimensional square lattice of size $L_x \times L_y$. A classical configuration of spins is denoted as $\mathcal{S}$, and we can associate a quantum state $|\mathcal{S}\rangle$ to each configuration. Any quantum state for this system can be written as a sum over these states with complex weights

$$|\Psi\rangle = \sum_{\{\mathcal{S}\}} \Psi(\mathcal{S}) |\mathcal{S}\rangle. \quad (1)$$

In the following, we will also consider subsets of degrees of freedom, for example horizontal rows on the lattice, or single sites, which we will denote as $\mathcal{S}_y$ and $\mathcal{S}_{x,y}$, respectively. In this paper, we will often sweep through the lattice row by row, going from left to right when sweeping through a specific row. To denote the partial state we obtained during a sweep up to and including some site with coordinates (x, y) , we use the notation $\mathcal{S}_{\leq x, \leq y}$. The same notation can be used on the level of rows; $\mathcal{S}_{\leq y}$ denotes the state of the first y rows.

A PEPS is defined by associating a tensor of order five to each point on the lattice, which can be visualized as

$$A_{\alpha\beta\gamma\delta}^s = \alpha \begin{array}{c} \beta \\ \text{---} \square \text{---} \gamma \\ \delta \end{array} s, \quad (2)$$

where the indices α , β , γ , and δ are the virtual indices and the index s denotes a physical degree of freedom, for example the up and down states of a spin degree of freedom. The PEPS tensor network is defined by contracting the virtual indices with those of the tensors associated to nearest neighbors on the lattice, visualized as

$$|\Psi(\{A\})\rangle = \begin{array}{c} \begin{array}{cccc} \square & \square & \square & \square \\ | & | & | & | \\ \square & \square & \square & \square \\ | & | & | & | \\ \square & \square & \square & \square \\ | & | & | & | \\ \square & \square & \square & \square \end{array} \end{array}, \quad (3)$$

where the connected lines imply a tensor contraction. In Eq. (3), the open virtual lines at the boundary are chosen to have dimension 1, and will be dropped hereafter.

In general, computing the norm or an expectation value for a PEPS cannot be done exactly with limited computational resources [47], and approximation techniques have to be considered. Representing the norm of a PEPS as

$$\langle \Psi(\{A\}) | \Psi(\{A\}) \rangle = \begin{array}{c} \begin{array}{cccc} \square & \square & \square & \square \\ | & | & | & | \\ \square & \square & \square & \square \\ | & | & | & | \\ \square & \square & \square & \square \\ | & | & | & | \\ \square & \square & \square & \square \end{array} \end{array}, \quad (4)$$

we can imagine that we can move through the lattice from the top to the bottom and progressively multiply the row-to-row transfer matrix. Since the dimension of the boundary operator scales exponentially with the number of columns, we can approximate it as a matrix product operator—a boundary MPO—with a tractable bond dimension χ_d . In each step of this row-to-row contraction [4],

$$\begin{array}{c} \begin{array}{cccc} \square & \square & \square & \square \\ | & | & | & | \\ \square & \square & \square & \square \\ | & | & | & | \\ \square & \square & \square & \square \end{array} \\ = \begin{array}{c} \square \text{---} \square \text{---} \square \text{---} \square \\ | \quad | \quad | \quad | \end{array}, \end{array} \quad (5)$$

the bond dimension of the boundary MPO grows. Therefore, in order to keep the bond dimension χ_d tractable, we approximate the boundary MPO by an MPO with a tractable bond

dimension

$$\approx \text{[Diagram of horizontal chain]} \quad (6)$$

This approximation step can be performed by local truncation steps or by solving a variational optimization problem, the complexity of which scales typically as $\mathcal{O}(\chi_d^2 D^6 + \chi_d^3 D^4)$.

In the spirit of the DMRG algorithm for MPS, the variational optimization of PEPS wave functions can be performed by a sweeping algorithm [4], where the PEPS tensors are updated consecutively until convergence. In the absence of a general prescription for bringing a PEPS into canonical form, the local optimization of a single tensor, however, suffers from an ill-conditioned inversion problem and is, therefore, unstable. Performing imaginary-time evolution with a Trotter-Suzuki splitting of the time-evolution operator [7,48] suffers from the same poor conditioning. Global optimization strategies are potentially more stable, but computing the gradient for the energy function is not straightforward. For infinite PEPS, gradient-based optimization strategies [9] have indeed shown to be more stable and yield states that are variationally optimal [9,10].

III. SAMPLING PEPS

Instead of considering the double layer and contracting it using a boundary MPO, another approach consists of interpreting the contraction over the physical indices of the PEPS in Eq. (4) as an explicit sum over the set of classical configurations

$$\langle \Psi(\{A\}) | \Psi(\{A\}) \rangle = \sum_{\{S\}} |\langle S | \Psi(\{A\}) \rangle|^2$$

$$= \sum_{\{S\}} \left(\text{[Diagram of grid with boundary vectors]} \right)^2 \quad (7)$$

where we define

$$\mathbb{1} = \text{[Diagram of red leg]} = \sum_{\{s\}} \text{[Diagram of circle with s]} \quad (8)$$

The number of terms in the sum in Eq. (7) scales exponentially as a function of the number of physical degrees of freedom. However, if the state is normalized, the terms in this sum are probabilities corresponding to the square amplitudes of the wave function and the sum in Eq. (7) corresponds to the normalization of this probability distribution. Furthermore, when considering expectation values of (sparse) operators

they can be rewritten as a classical expectation value over the probability distribution $|\Psi(S)|^2$

$$\begin{aligned} \langle \hat{O} \rangle &= \frac{\langle \Psi | \hat{O} | \Psi \rangle}{\langle \Psi | \Psi \rangle} \\ &= \sum_{\{S\}} \frac{|\Psi(S)|^2}{\langle \Psi | \Psi \rangle} \frac{\langle S | \hat{O} | \Psi \rangle}{\Psi(S)} \\ &= \sum_{\{S\}} \frac{|\Psi(S)|^2}{\langle \Psi | \Psi \rangle} O_{\text{loc}}(S), \end{aligned} \quad (9)$$

where we introduce the so-called local estimator of the operator $\hat{O}$

$$O_{\text{loc}}(S) = \langle S | \hat{O} | \Psi \rangle / \Psi(S) \quad (10)$$

and

$$p_{\Psi}(S) \equiv \frac{|\Psi(S)|^2}{\langle \Psi | \Psi \rangle} \quad (11)$$

can be interpreted as a probability distribution. Using sampling algorithms, this expectation value can be estimated by sampling this distribution and evaluating the mean of $O_{\text{loc}}(S)$ for all drawn samples

$$\langle \hat{O} \rangle \approx \frac{1}{N} \sum_{S \sim p_{\Psi}} O_{\text{loc}}(S). \quad (12)$$

The variance on the expectation value in Eq. (12) is then

$$\text{Var}_{p_{\Psi}}(\langle \hat{O} \rangle) = \left\langle \left(\frac{1}{N} \sum_{S \sim p_{\Psi}} (O_{\text{loc}}(S) - \langle \hat{O} \rangle) \right)^2 \right\rangle. \quad (13)$$

Note that in Eqs. (9) and (12), only amplitudes of the wave function appear

$$\Psi(S) = \text{[Diagram of 4x4 grid of blue tensors with circles labeled s on the legs]} \quad (14)$$

In contrast to the double-layer approach of Eq. (4), here we need to contract a *single-layer* PEPS. Similar to the double-layer contraction of Eq. (5), we can traverse the lattice row by row and multiply the transfer matrix with a boundary vector—in this case, a boundary MPS—which is kept at a tractable

bond dimension χ_s :

$$\begin{aligned}
 & \text{Diagram showing the contraction of single-layer PEPS tensors. The top part shows a chain of four green square tensors connected by horizontal lines, with each green tensor also connected to a blue square tensor below it. Each blue tensor has a small circle labeled 's' next to it. The bottom part shows the result of contracting the blue tensors, which are now represented as single horizontal lines between the green tensors, resulting in a simplified chain of four green tensors connected by single lines.} \\
 & (15)
 \end{aligned}$$

Contracting only single-layer PEPSs can be performed with a significant decrease in computational cost as compared to the double-layer contraction; in particular, it only scales as $\mathcal{O}(\chi_s^2 D^4 + \chi_s^3 D^2)$. Furthermore, it has been pointed out that χ_s can be chosen a lot smaller than χ_d [11]. However, finding a relation between the two, if at all present, is still an open question as these bond dimensions are strongly dependent on the model and the application. Yet, despite the scaling advantage for contracting a single layer PEPS, it has to be noted that this single layer contraction needs to be performed for many different samples S , as can be seen in Eq. (12). This is in contrast to the calculation of expectation values in a double-layer PEPS algorithm, where the expectation value can be expressed as a single contraction.

In the remainder of this section, we explain how to estimate observables such as in Eq. (9) using sampling. First, we review the Markov chain Monte Carlo method, which yields samples S distributed according to $p_\Psi(S)$. Afterwards we introduce an algorithm that yields samples S distributed according to a distribution $p_c(S)$, which lies close to the distribution $p_\Psi(S)$. Using importance sampling with the distribution $p_c(S)$, the observable in Eq. (9) can be estimated. Converse to MCMC, sampling from $p_c(S)$ can be performed directly and we introduce a way to systematically bring $p_c(S)$ closer to $p_\Psi(S)$, when desired.

A. Markov chain Monte Carlo

As is standard in the VMC approach, a straightforward way of sampling PEPS tensor networks is using a Markov chain Monte Carlo algorithm (MCMC). An example is the Metropolis Hastings algorithm [49,50], which provides a sampling procedure that requires few assumptions of the underlying probability distribution. In the Metropolis Hastings algorithm, starting from an initial configuration S^1 , one proposes a new configuration S^2 by perturbing S^1 according to a trial probability $\omega(S^1 \rightarrow S^2)$. The configuration S^2 is accepted with probability

$$P_{\text{acc}}(S^1 \rightarrow S^2) = \min\left(1, \frac{|\Psi(S^2)|^2 \omega(S^2 \rightarrow S^1)}{|\Psi(S^1)|^2 \omega(S^1 \rightarrow S^2)}\right). \quad (16)$$

The conditional probability to transition from state S^1 to S^2 is then given by

$$T(S^1 \rightarrow S^2) = \omega(S^1 \rightarrow S^2) P_{\text{acc}}(S^1 \rightarrow S^2). \quad (17)$$

Iterating this procedure yields a chain of configurations that produces samples from the distribution $|\Psi(S)|^2$ in the limit of

infinite number of samples. This is guaranteed by the detailed balance condition for the trial probabilities

$$|\Psi(S^1)|^2 T(S^1 \rightarrow S^2) = |\Psi(S^2)|^2 T(S^2 \rightarrow S^1), \quad (18)$$

which is in this case implied by the acceptance probabilities in Eq. (17). A clear disadvantage is that the transition between two configurations implies a degree of correlation between subsequent configurations. The choice of $\omega(S^1 \rightarrow S^2)$ determines the autocorrelation time and mixing dynamics, and hence the efficiency of the sampling algorithm. In VMC algorithms, the choice for $\omega(S^1 \rightarrow S^2)$ is typically a local update, involving one or two spins. This local update is performed consecutively in a sweeping fashion through the lattice. This choice yields large autocorrelation times, as many updates need to be performed to construct a sample S^f that is statistically independent from some initial sample S^i . Therefore, while performing the Markov chain, one only adds samples to the sample set periodically, with a time interval τ_s that has to be chosen greater than the autocorrelation time τ_a defined by the decay of correlations of the relevant operators at different times

$$\sum_{t=0}^{\infty} \langle O(S^{t+t'}) O(S^t) \rangle - \langle O(S^{t+t'}) \rangle \langle O(S^t) \rangle \sim \exp(-t'/\tau_a). \quad (19)$$

Another disadvantage of MCMC algorithms is that the Markov chain is started from an initial configuration S^i , that is typically not drawn according to the distribution $|\Psi(S)|^2$. Starting from S^i , consecutive samples are only distributed according to $|\Psi(S)|^2$ in the limit of infinite time. In practice, one performs an equilibration period τ_i that is sufficiently long such that the value of an observable becomes approximately stationary when estimated using different Markov chains. The period τ_i is determined by the gap between the largest and second largest eigenvalues of $T(S^1 \rightarrow S^2)$, and the samples drawn during this equilibration period are not usable for estimating expectation values. Determining τ_i and τ_s is computationally hard and, as such, heuristics are used to determine them in practice.

The choice for local updates in VMC algorithms is imposed by the fact that during every optimization step, a different probability distribution $|\Psi(S)|^2$ is encountered. Therefore, it is impractical to determine a rule $\omega(S^1 \rightarrow S^2)$ that is optimal at every optimization step, as this requires much knowledge about the structure of the probability distribution. When this knowledge does exist—e.g., for classical Boltzmann distributions with nearest-neighbor interactions—highly efficient trial probabilities can be constructed. An example is the Wolff algorithm [51] performing updates of whole clusters.

B. Direct sampling

In order to alleviate these shortcomings, we introduce a method for drawing independent samples from a PEPS tensor network, also called direct sampling (DS). In order to have an efficient direct-sampling scheme, we (i) directly sample from an effective distribution $p_c(S)$ that is not too far from $p_\Psi(S)$, (ii) provide a way to compute the ratio $|\Psi(S)|^2/p_c(S)$ for each sample that we draw and (iii) use this ratio in an importance

sampling algorithm to determine expectation values. The distribution p_c will be constructed as a sequence of conditional probabilities for each row.

As explained in Sec. II, we can view the contraction of the double-layer network as the evolution of a one-dimensional density matrix. We define

$$\rho_1 = \text{[Diagram: A horizontal chain of four blue square tensors connected by vertical lines. The left and right ends have open bonds.]}, \quad (20)$$

and

$$\rho_2 = \text{[Diagram: A 2x4 grid of blue square tensors. The top row is connected horizontally, and the bottom row is connected horizontally. Vertical lines connect corresponding tensors in the two rows. The left and right ends have open bonds.]}, \quad (21)$$

where the uncontracted bonds denote a density matrix defined on the vector space of the boundary. The evolution of the density matrix can be written as

$$\rho_y = \sum_{\{S_y\}} c_y(S_y) \rho_{y-1} c_y(S_y)^\dagger, \quad (22)$$

where $c_y(S_y)$ is the y -th row of the PEPS projected on the physical basis states S_y . This equation can be seen as a discrete version of the Lindblad equation

$$\frac{d\rho}{dt} = -i(H_{\text{eff}}\rho - \rho H_{\text{eff}}^\dagger) + \sum_m c_m \rho c_m^\dagger, \quad (23)$$

where we identify the effective Hamiltonian $H_{\text{eff}} = -\frac{i}{2}\mathbb{1}$ and the Lindblad operators $c_m = c_y(S_y)$. One approach for integrating the Lindblad equation consists of stochastically obtaining a set of trajectories and estimating expectation values from this set. The trajectories consist of time evolutions of pure states, where the dynamics according to H_{eff} is alternated with applications of the operators c_m , the so-called quantum jumps. The full algorithm is detailed in Ref. [52] and an overview can be found in Appendix A.

For our purpose of sampling PEPS double-layer networks, we can use this quantum jump algorithm to simulate trajectories that can be identified by configurations $\mathcal{S}$. We start from the trivial boundary density operator with dimension one

$$\rho_0 = \mathbb{1} = \text{[Diagram: Four separate vertical lines, each representing a dimension-1 state.]}. \quad (24)$$

Locally, this density operator can be written as the outer product of a one-dimensional unit vector with itself. For the full boundary, this yields $\rho_0 = |\phi_0\rangle_b \langle\phi_0|_b$, with

$$|\phi_0\rangle_b = \text{[Diagram: Four vertical lines, each with a dot at the top, representing a unit vector in the boundary space.]}, \quad (25)$$

where the subscript denotes that the bra and ket pertain to the vector space of the boundary, and not the physical space. The boundary vector $|\phi_0\rangle_b$ can be written in MPS form with bond dimension one

$$|\phi_0\rangle_b = \text{[Diagram: A horizontal line with four dots, representing an MPS with bond dimension one.]}. \quad (26)$$

Starting from this trivial boundary MPS, we propose a sequence of three steps for moving down one row y in the single-layer PEPS, finding a sample S_y for the row's configuration space as we go.

(1) Apply a row of the single-layer PEPS tensors to the boundary MPS from the previous step,

$$\mathbf{c}_y |\phi_{y-1}\rangle_b = \text{[Diagram: A horizontal chain of four blue square tensors with an additional index (dot) on top of each, connected by vertical lines to a horizontal line above.]}, \quad (27)$$

$$= \text{[Diagram: A horizontal chain of four orange square tensors connected by vertical lines.]},$$

where we introduce the tensor $\mathbf{c}_y$ as the collection of all operators $c_y(S_y)$, i.e., it is an MPS-like state, where every tensor has an extra index originating from the physical degrees of freedom. After obtaining $\mathbf{c}_y |\phi_{y-1}\rangle_b$, the bond dimension of the boundary has grown, and it should be truncated if necessary. In general, we will impose a maximal bond dimension on $\mathbf{c}_y |\phi_{y-1}\rangle_b$, which we will denote χ_s .

(2) Sample one of the configurations S_y according to the probability distribution

$$p_c(S_y | \mathcal{S}_{<y}) = \frac{\langle \phi_{y-1} |_b c(S_y)^\dagger c(S_y) | \phi_{y-1} \rangle_b}{\sum_{\{S_y\}} \langle \phi_{y-1} |_b c(S_y)^\dagger c(S_y) | \phi_{y-1} \rangle_b}. \quad (28)$$

This sampling step is equivalent to the sampling of an MPS and can be efficiently performed by using the perfect-sampling approach for unitary tensor networks [45]. For convenience, we have summarized this algorithm in Appendix B.

(3) When a configuration $S_y = (S_{1,y}, \dots, S_{L_y,y})$ is found, one obtains a new boundary MPS $|\phi_y\rangle_b$ by projecting the physical degrees of freedom in $\mathbf{c}_y |\phi_{y-1}\rangle_b$ onto the configuration S_y ,

$$|\phi_y\rangle_b = \text{[Diagram: A horizontal chain of four orange square tensors with circles above them labeled S_{1,y}, S_{2,y}, S_{3,y}, S_{4,y}, connected by vertical lines to a horizontal line above.]}, \quad (29)$$

$$= \text{[Diagram: A horizontal line with four dots, representing the resulting boundary MPS.]}. \quad (29)$$

For convenience, we divide this boundary MPS by $\sqrt{p_c(S_y | \mathcal{S}_{<y})}$.

After going through all L_y rows of the PEPS, we obtain a sample configuration $\mathcal{S}$, sampled with the sequence of conditional probabilities

$$p_c(\mathcal{S}) = p_c(S_1) p_c(S_2 | S_1) \dots p_c(S_{L_y} | \mathcal{S}_{<L_y}). \quad (30)$$

This equation highlights the viewpoint that $p_c(\mathcal{S})$ is built as a sequence of conditional probability distributions of row configurations. As such, the algorithm follows the same strategy as autoregressive sampling techniques and also the direct sampling algorithm for one dimensional loop-free tensor networks [38,39,45]. In these algorithms, one assumes that marginal probability distributions can be efficiently calculated or have a specific structure by construction of the model. Note that we arrived at Eq. (30) by assuming that the marginals of $p_\Psi(\mathcal{S})$ are equal to the unit reduced density matrix. Hence, $p_c(\mathcal{S})$ is

related but not equal to $p_\Psi(\mathcal{S})$. Indeed, the prescription for p_c in Eq. (28) is a drastic approximation for the true marginal probability distributions of $p_\Psi(\mathcal{S})$ for the variables $\mathcal{S}_y$, as it ignores all other rows in the network. We can systematically improve on the above scheme such that the distribution $p_c(\mathcal{S})$ comes closer to $p_\Psi(\mathcal{S})$. Indeed, we can improve the approximation in step 2 of the algorithm by tracing out the environment of the trailing rows approximately,

$$E_m(\mathcal{S}_{>y}) = \text{[Diagram of double-layer contraction]} \quad (31)$$

These environments are MPOs, that can be truncated to a maximal bond dimension χ_m . In the sampling algorithm, using these environments boils down to inserting this MPO in the above prescription for the marginal distribution p_c

$$p_c^m(\mathcal{S}_y | \mathcal{S}_{<y}) = \frac{\langle \phi_{y-1} | b c(\mathcal{S}_y)^\dagger E_m(\mathcal{S}_{>y}) c(\mathcal{S}_y) | \phi_{y-1} \rangle_b}{\sum_{\{\mathcal{S}_y\}} \langle \phi_{y-1} | b c(\mathcal{S}_y)^\dagger E_m(\mathcal{S}_{>y}) c(\mathcal{S}_y) | \phi_{y-1} \rangle_b}. \quad (32)$$

By increasing the environment bond dimension χ_m , the distribution $p_c^m(\mathcal{S})$ that is sampled in our scheme will converge to $p_\Psi(\mathcal{S})$. Finding the environments E_m , however, requires a boundary-MPO contraction of the double-layer network, which we were set out to avoid in the first place. We will show in the following section that using small values of χ_m yield environments that are computationally cheap to compute, compared to the rest of the algorithm, and already yield satisfactory results—although a full double-layer computation with this value of the boundary-MPO bond dimension would not yield an acceptable accuracy.

The expectation value in Eq. (9) can now be written as

$$\begin{aligned} \langle \hat{O} \rangle &= \frac{1}{\langle \Psi | \Psi \rangle} \sum_{\{\mathcal{S}\}} p_c(\mathcal{S}) \frac{|\Psi(\mathcal{S})|^2}{p_c(\mathcal{S})} O_{\text{loc}}(\mathcal{S}) \\ &\approx \frac{1}{N \langle \Psi | \Psi \rangle} \sum_{\mathcal{S} \sim p_c} \frac{|\Psi(\mathcal{S})|^2}{p_c(\mathcal{S})} O_{\text{loc}}(\mathcal{S}). \end{aligned} \quad (33)$$

In this equation, the term $|\Psi(\mathcal{S})|^2 / p_c(\mathcal{S})$ can be readily computed after a sample is drawn by calculating the norm of Eq. (29) as $|\phi_y\rangle_b$ is essentially a single-layer contraction $\Psi(\mathcal{S})$ divided by the root of $p_c(\mathcal{S})$ as defined in Eq. (30). The norm $\langle \Psi | \Psi \rangle$ can also be computed by using the resulting samples as it can be written as

$$\langle \Psi | \Psi \rangle = \sum_{\{\mathcal{S}\}} p_c(\mathcal{S}) \frac{|\Psi(\mathcal{S})|^2}{p_c(\mathcal{S})} \approx \frac{1}{N} \sum_{\mathcal{S} \sim p_c} \frac{|\Psi(\mathcal{S})|^2}{p_c(\mathcal{S})}. \quad (34)$$

Equation (33) is expressed as an importance sampling procedure, sampling from $p_c(\mathcal{S})$. Hence, the estimate of the expectation value is asymptotically unbiased. The variance on

the expectation value in Eq. (33) is defined as

$$\text{Var}_{p_c}(\langle \hat{O} \rangle) = \left\langle \left(\frac{1}{N} \sum_{\mathcal{S} \sim p_c} \left(\frac{|\Psi(\mathcal{S})|^2 O_{\text{loc}}(\mathcal{S})}{\langle \Psi | \Psi \rangle p_c(\mathcal{S})} - \langle \hat{O} \rangle \right) \right)^2 \right\rangle, \quad (35)$$

i.e., the expected square deviation of the estimate of the observable using N samples from the true observable.

It is clear that the disadvantages related to autocorrelation in a Markov chain do not arise in this algorithm. The samples are drawn independently from each other and no equilibration period is needed. A second advantage is the possibility to estimate the norm of a PEPS as shown in Eq. (34), which is not possible when sampling from p_Ψ by a Markov chain. Furthermore, as the algorithm has no interdependence between samples, it can be executed in a massively parallel fashion.

C. Independent samples as Markov chain transitions

Finally we can combine both sampling approaches. Indeed, we can take the samples drawn according to p_c as the transition probabilities in a Markov chain. Concretely, we take as transition probabilities $\omega(\mathcal{S}^1 \rightarrow \mathcal{S}^2) \equiv \omega(\mathcal{S}^2) = p_c(\mathcal{S}^2)$. In this case, the proposed configuration is independent from the previous one, and the new configuration is accepted with the acceptance probability defined in Eq. (16). As detailed balance [Eq. (18)] is satisfied, the resulting distribution from which the samples are drawn is p_Ψ . This algorithm still has the advantage of drawing independent samples and being massively parallelizable. The downside is that an equilibration period is needed because, at the start of this Markov chain, the samples are drawn according to p_c , which is not equal to p_Ψ . This approach is similar to that of Ref. [46], where the transition probabilities $T(\mathcal{S}^2)$ were determined by approximately contracting a single layer tensor network representing a classical Boltzmann distribution.

IV. STOCHASTIC OPTIMIZATION OF PEPS

Given that we can efficiently sample a PEPS wave function and compute expectation values, we now give details on how to variationally optimize the PEPS tensors with respect to a given model Hamiltonian.

The first issue concerns the initialization of the PEPS tensors. In recent works, an initial state was constructed by running the simple-update algorithm for PEPS, stating that starting from an initial PEPS with randomly initialized parameters leads to poor statistics and local minima [33]. In this work, however, we have observed that this additional step is unnecessary by choosing a PEPS of which the weights are randomly initialized from a uniform distribution between zero and one. With this initialization, the distribution of the expansion coefficients $\Psi(\mathcal{S})$ is narrowly distributed around a nonzero value. This is qualitatively dissimilar from the distribution obtained with a PEPS whose weights are initialized with positive and negative weights. Indeed, as a single-layer PEPS contraction can be viewed as a sum of products of tensor elements, this sum will have many canceling contributions, leading to a distribution centered and peaked around zero. When considering local operators $O_{\text{loc}}(\mathcal{S})$, the expansion

coefficients enter in the denominator and expansion coefficients close to zero lead to large contributions to the local operator. Note that, with a finite sample set, the large contributions lead to poor statistics of the expectation value. On the other hand, when the expansion coefficients are narrowly distributed around a nonzero value, the matrix elements of the operator $\langle S|\hat{O}|S'\rangle$ have a larger effect on the value of the local operator and the expectation value has low variance. We also rescale the weights by properly normalizing the PEPS, ensuring that the mean of the distribution $|\Psi(S)|^2/p_c(S)$ is equal to one.

It is often worthwhile to consider symmetries of the models under study when optimizing variational ground states. In this work, when present, we will consider eigenstates of the U(1) symmetry group by projecting the PEPS on the subspace with total spin projection equal to a fixed value Q

$$|\Psi_{S^z=Q}\rangle = \hat{P}_{S^z=Q}|\Psi\rangle, \quad S^z = \sum_{x,y} S_{x,y}^z. \quad (36)$$

This projector can be efficiently implemented by sampling only configurations S with the constraint $\langle S|\hat{S}_z|S\rangle = Q$. Indeed, we start the sampling as is described above and we keep track of the total U(1) charge accumulated during the sampling. At some point, we reach a bound that is determined by the desired value Q and the total number of spins; from that point on, we define the distributions $p_c(S_{x,y}|S_{<x,<y})$ such that it is zero for $S_{x,y}$ equal to the species that reached its bound. In this way, we prohibit the amount of spins of a species to grow beyond its bound.

Note that this approach for imposing a U(1) symmetry (or any Abelian symmetry) is different than the standard one in the tensor-network formalism, where a global (in general, non-Abelian) symmetry of a given tensor-network state can be encoded as a constraint on the level of the local tensors. Here, we apply a global projector on the full state by constraining the sampling procedure. In a direct contraction, applying this projector would require a prohibitively large bond dimension.

After initialization, the parameters of the PEPS are updated using an estimate of the gradient of the expectation value of the Hamiltonian with respect to the tensors defined in Eq. (2). The expression for these gradients is

$$g = \frac{\partial \langle \hat{H} \rangle}{\partial A_{\alpha\beta\gamma\delta}^{(x,y)}(s)} = \sum_{\{S\}} \frac{|\Psi(S)|^2}{\langle \Psi|\Psi \rangle} \frac{\partial \ln(\Psi^*(S))}{\partial A_{\alpha\beta\gamma\delta}^{(x,y)}(s)} (H_{\text{loc}}(S) - \langle \hat{H} \rangle), \quad (37)$$

where we denote the coordinates of the different tensors as (x, y) . The gradient provides the direction which maximizes the expectation value of the Hamiltonian locally in parameter space; by updating the parameters in the opposite direction using a small step size, we move the PEPS state closer to the minimum of the energy, which is known as the stochastic gradient-descent (SGD) algorithm. Note that the gradient as calculated by sampling, and more generally any expectation value, is inherently noisy due to the finite amount of samples. The noise will have an effect on the gradient direction, i.e., not pointing exactly in the opposite direction of steepest descent. However, the accumulation of different noisy steps will yield a trajectory in parameter space that moves in the direction of lower energies because the average direction of different noisy

gradients is still close to the direction of steepest descent, and thus a descent direction. Furthermore, the level of noise can be controlled by taking more samples into account in the estimation of expectation values.

To alleviate this problem, one can use information from previous steps. An example is the ADAM optimization algorithm [53], that uses an exponentially weighted average of previous gradients to construct the direction of the current step. This is combined with a parameter-dependent step size, scaling as the exponentially weighted average of the norm of the gradient of a parameter.

Furthermore, we can refine the gradient descent approach by including information about the shape of the manifold described by the PEPS ansatz and proposing an update direction based on this information combined with the gradient. In the VMC context, this is the so-called stochastic reconfiguration (SR) approach [54,55], also known as natural gradient descent [56]. An alternative viewpoint is that it includes information about the metric of the PEPS manifold or, equivalently, performs an imaginary time evolution step on the PEPS manifold according to the time-dependent variational principle [57–59]. In any case, we find a more natural descent direction by solving the following linear equation for the direction $\mathbf{x}$

$$\mathbf{S}\mathbf{x} = \mathbf{g}, \quad (38)$$

where we define the matrix $\mathbf{S}$ as

$$\mathbf{S}_{i,j} = \left\langle \frac{\partial \ln(\Psi^*)}{\partial p_i} \frac{\partial \ln(\Psi)}{\partial p_j} \right\rangle - \left\langle \frac{\partial \ln(\Psi^*)}{\partial p_i} \right\rangle \left\langle \frac{\partial \ln(\Psi)}{\partial p_j} \right\rangle. \quad (39)$$

Here the indices i and j denote collective indices for all variational parameters in the PEPS wave function. For a given sample set, the matrix $\mathbf{S}$ can be singular and hence the inversion needed to solve Eq. (38) for $\mathbf{x}$ can fail. In practical applications described below, we will regularize the matrix $\mathbf{S}$ by shifting the diagonal of the matrix with an amount δ . In all simulations below, we use $\delta = 0.01$. Solving the linear system of equations in Eq. (38) can be done with an iterative solver such as the conjugate gradient method. Using this algorithm, one has to perform the multiplication of $\mathbf{S}$ with some arbitrary vector, which can be performed efficiently by sampling $\mathbf{S}$, scaling with the number of variational parameters, i.e., $\mathcal{O}(D^4)$.

V. RESULTS

A. Performance benchmark

As noted in Sec. III B, direct sampling aims to solve downsides of MCMC algorithms such as the problem of autocorrelation and equilibration. The problem of equilibration translates to the fact that MCMC algorithms can not be parallelized indefinitely, because the sampling during the equilibration period is inherently serial. In Fig. 1, we show the wall time as a function of the number of CPU cores for drawing a sample set of $N = 1000$ samples from a 10×10 PEPS with bond dimension $D = 4$. We compare direct sampling with MCMC, both using a cutoff $\chi_s = 4$ in the contractions. For MCMC, we use an equilibration period of $\tau_i = 100$ sweeps and $\tau_s = 1$ sweep between consecutive samples. We define a sweep as performing a local update once on every site of the lattice, visiting the lattice sites row by row, going from left to right within every row. For direct sampling, we

FIG. 1. Comparison of wall time as a function of the inverse of the number of CPU cores used for direct sampling and MCMC. Every point is the median of ten separate runs. Every run consists of drawing $N = 1000$ samples, where an equilibration time of $\tau_i = 100$ sweeps was used for the MCMC runs. The inset shows the relative speedup of the direct sampling procedure compared to MCMC for constructing a full set of samples.

use a cutoff $\chi_m = 4$ for the environments of the marginals. In theory, the wall time of direct sampling scales as $t = \mu/N_c$ and that of MCMC scales as $t = \nu\tau_i + \xi\tau_s/N_c$, where N_c is the number of cores used. In practice, computational overhead increases the runtime on an actual computer. It is clear that direct sampling is twice as fast for a low number of CPU cores, while it progressively becomes more efficient when using more CPU cores, reaching a five-fold speedup for 20 CPU cores. The MCMC algorithm suffers from the combination of numerical overhead and the equilibration when using more than 10 CPU cores, while overhead is much less severe in the case of direct sampling.

B. Benchmark on classical spin models

As a proof of principle, we apply our sampling algorithm to drawing samples according to the partition function of the classical two-dimensional Ising model. The classical Ising model is defined by the energy function

$$E(S) = -J \sum_{\langle i,j \rangle} S_{i,j} S_{k,l}, \quad (40)$$

where $S_{i,j}$ are classical spin variables and the sum iterates over all pairs of nearest neighbors. The partition function can be written as the norm of a PEPS wave function [60]. Indeed, by defining the tensors

$$A_{\alpha\beta\gamma\delta}(s) = \begin{cases} 1, & \text{if } \alpha = \beta = \gamma = \delta = s \\ 0, & \text{otherwise} \end{cases} \quad (41)$$

and

$$B = \begin{bmatrix} e^{1/2kT} & e^{-1/2kT} \\ e^{-1/2kT} & e^{1/2kT} \end{bmatrix}, \quad (42)$$

and arranging the tensors A on the lattice sites of the square lattice and the tensors B on the links between the lattice sites, one sees that the resulting PEPS state represents the vector of square root probabilities of the Ising model. In the following,

FIG. 2. Projection of data points using different sampling approaches on two principal components, for a PEPS encoding the classical Ising model. The principal components are those with largest explained variance of the correlation matrix of the sample set constructed with direct sampling. The first component coincides with magnetization of the samples. MCMC, which is in this case equivalent to the Metropolis Hastings algorithm applied to the Boltzmann distribution, is clearly unable to sample both modes of the probability distribution. The Wolff algorithm and the direct sampling approaches do not suffer from this problem.

we compare our sampling algorithm with the Metropolis Hastings algorithm using single spin flips, and the Wolff algorithm [51]. The former is a very versatile variant of MCMC, which is also typically used for the optimization of quantum states. The latter is specifically designed for Ising-like models, and is well-known for its favorable scaling of the (integrated) autocorrelation time. In Fig. 2, we show a principal component analysis of sampled configurations of a 10×10 Ising model in the ferromagnetic phase. Specifically, we show the two components with largest explained variance, i.e., the projection of the data on the two eigenvectors with largest eigenvalue of the covariance matrix of the data [61]. For every algorithm, we sample 10 000 configurations and subsample with equal intervals 1000 samples of our PEPS sampling algorithm and 100 samples of the other algorithms. The first principal component is closely related to the magnetization, and shows the bimodal distribution of it. Our sampling algorithm is symmetric with respect to the sign of the magnetization, unlike the MCMC algorithm which tends to be unable to cross the large energy barrier between the two modes. Also when using our sampling algorithm as transitions in a Markov chain (denoted DS+MC), the resulting samples have equal weight in both modes of the magnetization, as the transitions are independent of the previous samples. This shows that our sampling algorithm can be used to cross large barriers in probability space, as is necessary in, e.g. glassy models [62–64]. As a second benchmark, we turn to the autocorrelation of observables in a Markov chain. The autocorrelation of the estimator of an observable $\hat{O}$ is defined as

$$\Gamma_{\hat{O}}(t) = \frac{1}{N} \sum_{t'=1}^N O(t') O(t' + t) - \langle \hat{O} \rangle^2, \quad (43)$$

FIG. 3. Comparison of sampling efficiency of a 10×10 classical square-lattice Ising model between MCMC, Wolff and our direct sampling + MC approach. For the direct sampling + MC approach, we use $\chi_m = D = 2$. (a) The integrated autocorrelation time as defined in Eq. (44) as a function of J/kT . MCMC shows clear slowing down in the vicinity of the critical point, while this effect is much less severe in the case of direct sampling + MC and Wolff. (b) Acceptance rate of the Markov chain used in MCMC and direct sampling + MC.

where $O(t)$ is the estimator of the operator $\hat{O}$ at time t in the chain. The integrated autocorrelation time for this observable is then defined as

$$\tau_{\hat{O}}^{\text{int}} = 1 + 2 \sum_{t=1}^{\infty} \frac{\Gamma_{\hat{O}}(t)}{\Gamma_{\hat{O}}(0)}. \quad (44)$$

As no Markov chain is employed in our direct sampling algorithm, the integrated autocorrelation time is equal to one. For algorithms where a Markov chain does appear, we estimate the integrated autocorrelation time in Fig. 3(a). We perform ten Markov chains where we measure the absolute magnetization after every update step. For the direct sampling combined with Markov chain and Wolff algorithms, we take a chain length of $N = 10^4$ and for the MCMC algorithm we perform local updates and put a sample in the set after every sweep, i.e., 10^6 local updates result in $N = 10^4$ samples. The autocorrelation function is averaged over the ten distinct chains. We then calculate the integrated autocorrelation time by performing the sum up to a cutoff which we take as the time at which the autocorrelation function drops below a certain value, i.e., $\log(\Gamma_{\hat{O}}(t)) < -3$. Beyond that point, the correlation time becomes very noisy and the sum does not

converge. The portion that is left out by doing this procedure can be estimated to be an exponentially decaying function. Taking this exponentially decaying function in account leads to an underestimation of the integrated autocorrelation time by approximately 5%. The MCMC algorithm has the largest autocorrelation time due to its local updates and relatively low acceptance ratio. As indicated in Fig. 3(b) our direct sampling + MC algorithm has higher acceptance ratio and global updates, the autocorrelation time is significantly lower. The autocorrelation time has an effect on the accuracy of estimates of observables, as it increases the variance of expectation values as determined by Eq. (13). Specifically, it can be proven that the variance on an observable estimated using Markov chain sampling is [65]

$$\begin{aligned} \text{Var}_{p_{\psi}, \text{MCMC}}(\langle \hat{O} \rangle) &= \frac{\tau_{\hat{O}}^{\text{int}}}{N} \text{Var}_{p_{\psi}}(O_{\text{loc}}) \\ &= \frac{1}{N_{\text{eff}, \text{MCMC}}} \text{Var}_{p_{\psi}}(O_{\text{loc}}), \end{aligned} \quad (45)$$

where we introduced the effective sample set size $N_{\text{eff}, \text{MCMC}} = N/\tau_{\hat{O}}^{\text{int}}$. The variance using a MCMC algorithm is thus $\tau_{\hat{O}}^{\text{int}}$ times higher than that obtained using a direct sampling algorithm of p_{ψ} . Conversely, in our direct sampling algorithm, we draw samples from a distribution p_c , and the importance sampling correction in the observable estimation of Eq. (33) has a similar effect on the variance on that observable, defined in Eq. (35). One can prove that the following relation exists [66,67]:

$$\begin{aligned} \text{Var}_{p_c}(\langle \hat{O} \rangle) &= \frac{\frac{1}{N} \sum_{S \sim p_c} (|\Psi(S)|^2/p_c(S))^2}{N} \text{Var}_{p_{\psi}}(O_{\text{loc}}) \\ &= \frac{1}{N_{\text{eff}, \text{DS}}} \text{Var}_{p_{\psi}}(O_{\text{loc}}), \end{aligned} \quad (46)$$

where we introduced the effective sample set size $N_{\text{eff}, \text{DS}} = N^2 / \sum_{S \sim p_c} (|\Psi(S)|^2/p_c(S))^2$. From the above equation, it is clear that the variance of observables obtained using our direct sampling approach depends on the fluctuations in $|\Psi(S)|^2/p_c(S)$.

Both MCMC and direct sampling have suboptimal variances due to the effective sample set sizes $N_{\text{eff}, \text{MCMC}}$ and $N_{\text{eff}, \text{DS}}$. To show their differences, we compare the variance on the absolute magnetization of a 10×10 classical Ising model for different number of samples N using different sampling schemes in Fig. 4. All curves follow the expected scaling as $1/N$. The vertical offset is determined by the effective sample set size, either the effect of $\tau_{\hat{O}}^{\text{int}}$ or the fluctuations in $|\Psi(S)|^2/p_c(S)$. For the Markov chain based sampling schemes, the cluster updates in the Wolff algorithm and the accompanying low autocorrelation times yield the lowest variances. Using our direct sampling scheme with $\chi_m = 2$ yields little fluctuations in $|\Psi(S)|^2/p_c(S)$ and the effective sample set size is large, yielding even lower variances than the Wolff algorithm.

C. Effect of the marginal environment

As explained in Sec. III B, the distribution p_c approaches the distribution p_{ψ} systematically when increasing the cutoff χ_m of the environments of the marginal probability

FIG. 4. Variance on the absolute magnetization of a 10×10 classical square-lattice Ising model as a function of sample set size, for different sampling algorithms. The samples were drawn from the Boltzmann distribution at $kT/J = 1.9$. All algorithms follow the expected $1/N$ behavior. MCMC has the largest variance, due to the large autocorrelation time. The Wolff and direct sampling + MC algorithms have similar variances. The direct sampling algorithm without MC step has lowest variances.

distributions. In Fig. 5, we show histograms of the ratio of the probabilities p_c and p_Ψ for different cutoffs χ_m . We draw 1000 samples using our direct sampling approach from a 10×10 PEPS with $D = 4$, optimized on the ground state of the anti-ferromagnetic Heisenberg model (AFH) with Hamiltonian

$$\hat{H} = \sum_{\langle i,j \rangle} \hat{S}_i^x \hat{S}_j^x + \hat{S}_i^y \hat{S}_j^y + \hat{S}_i^z \hat{S}_j^z, \quad (47)$$

where $\hat{S}_i^{x,y,z}$ is the spin-1/2 operator on spin i in the x, y, z direction. We use the notation $\chi_m = 0$ to indicate that no environment was used, i.e., the unit MPO. Direct sampling with $\chi_m = 0$ leads to a large mismatch between p_c and p_Ψ . Using $\chi_m = 1$, i.e., an environment that is still spatially uncorrelated but allowing nontrivial correlations between the virtual dimensions of the PEPS tensors, already improves the distance between both distributions. For $\chi_m = 4$, the histograms start

FIG. 5. Histograms of $p_c(S)/p_\Psi(S)$ for 1000 samples drawn from p_c , for different cutoffs χ_m of the environments of the marginals. The inset shows the Kullback-Leibler divergence between p_c and p_Ψ .

FIG. 6. Similarity between p_c and p_Ψ . (a) Kullback-Leibler divergence $D_{KL}(p_c||p_\Psi)$ as a function of the cutoff of the environments of the marginal probability distributions χ_m for PEPS wave functions optimized on the ground state of the AFH with different bond dimensions D . The Kullback-Leibler divergence drops sharply for low values of χ_m , signaling that the overlap between p_c and p_Ψ improves greatly with increasing χ_m . (b) effective sample set size $N_{\text{eff,DS}}$ as defined in Eq. (46), as a function of χ_m . The black line indicates the number of samples drawn, i.e., $N = 1000$. The effective sample set sizes reach values close to this optimum for $\chi_m \approx D$ and keep improving when increasing χ_m .

to peak sharply around $p_c/p_\Psi = 1$. In the inset, we show the Kullback-Leibler divergence between p_c and p_Ψ , defined as

$$D_{KL}(p_c||p_\Psi) = \sum_{\{S\}} p_c(S) \log \left(\frac{p_c(S)}{p_\Psi(S)} \right), \quad (48)$$

which quantifies the relative entropy between the distributions p_c and p_Ψ . In Fig. 6(a), we show $D_{KL}(p_c||p_\Psi)$ as a function of χ_m for PEPS wave functions with different bond dimensions D . For values lower than $\chi_m \approx D$, the Kullback-Leibler divergence drops sharply, dropping up to three orders of magnitude compared to $\chi_m = 0$. In this region, the distribution p_c/p_Ψ goes from a broad distribution to a narrow distribution centered around $p_c/p_\Psi = 1$, as can be seen in Fig. 5. In Fig. 6(b), we show the effective sample set size $N_{\text{eff,DS}}$ as a function of χ_m . When $\chi_m \approx D$, $N_{\text{eff,DS}}$ lies close to the number of samples drawn, indicating that variances obtained by sampling directly from p_c are equally low as variances obtained by sampling from p_Ψ . Note that, although a double layer contraction, the environments of the marginals can be calculated in advance and constitute a negligible cost compared to the cost of sampling for all shown χ_m . More generally, for the values of $\chi_m \approx D$ used here, the calculation of the marginals is $\mathcal{O}(D^2)$ more expensive than drawing a single sample. For bond

FIG. 7. Comparison between boundary MPS combined with direct sampling, and boundary MPO methods for calculating observables. The observables are calculated for a $D = 4$ ground state of a 11×11 Transverse Field Ising model at $g = 3.2$. For the sampled observables, $N = 10^6$ samples were used to largely eliminate statistical fluctuations. (a) Square norm per site of the PEPS. The boundary MPS method converges more quickly than the boundary MPO method as a function of the cutoff. (b) Spin in transverse field direction of the center site. The error in the environment of the boundary MPO method is largely canceled out as the environment appears in the numerator and the denominator of the expectation value. Hence, the error on the expectation value converges more rapidly than the expectation value of the norm in (a).

dimensions $D = 1$ – 10 used in this work, this cost is lower than constructing all samples in the sample set, which is typically of size $N = 1000$ – $10\,000$.

D. Boundary MPS accuracy

According to Eqs. (13) and (35), sampling of observables introduces a variance on the resulting estimates. Apart from this variance, the algorithm also introduces a bias, due to the cutoff χ_s of the single-layer boundary MPS. This bias is also present in double-layer contractions, where it originates from the cutoff χ_d of the boundary MPO. The bias can be systematically reduced by increasing χ_s or χ_d . The systematic reduction for both algorithms is shown in Fig. 7. In Fig. 7(a), we show the norm per site of the PEPS as a function of χ_s and χ_d . It is clear that the bias on the norm is small for the boundary MPS method, and converges rapidly to zero at small χ_s . The bias is markedly larger for the boundary

FIG. 8. Energy as a function of iteration steps for a 10×10 PEPS with $D = 4$ using stochastic reconfiguration, gradient descent, and ADAM. For all runs, expectation values were obtained with $N = 1000$ samples. All runs used a fixed step size, i.e., that for which convergence was fastest. This is 0.05 for SR and 0.01 for SGD and ADAM. Note that SR requires an additional matrix inversion compared to SGD and ADAM. This step scales as $\mathcal{O}(D^4)$, i.e., more favourable than the cost of single-layer contractions. In this case specifically, the matrix inversion adds roughly 30% of runtime per iteration step.

MPO method, and converges more slowly to zero. However, in typical expectation values containing local operators, the same environment enters the numerator and the denominator, hence largely canceling the errors on the environment. This can be observed in Fig. 7(b), where the expectation value of the spin in the direction of the transverse field is shown as a function of χ_s or χ_d . The convergence of this expectation value for the boundary MPO algorithm is much faster than the convergence of the norm, although the norm is part of the calculation of this expectation value. This effect is also present for the boundary MPS algorithm, but is to a lesser extent important as the unnormalized expectation values already converge for low χ_s . Comparing the convergence of the norm in Fig. 7(a) as a function of χ_d , with the convergence of the effective sample set size $N_{\text{eff,DS}}$ in Fig. 6(b) as a function of χ_m , it is clear that χ_m can be taken much smaller than χ_d , although they are both double layer contractions.

All truncations were performed using singular value decompositions. Based on these results and systematic investigations in Refs. [11,33], we use the heuristic $\chi_s = 3D$ in all numerical experiments in this paper. Figure 7 shows that this is clearly sufficient.

E. Optimizing ground states of quantum spin models

In this section, we optimize a PEPS variationally using samples drawn according to our direct sampling algorithms.

1. Stochastic reconfiguration versus gradient descent

In Fig. 8, we show how the energy decreases as a function of the number of optimization steps for the stochastic gradient descent update [Eq. (37)], ADAM update, and the stochastic reconfiguration update [Eq. (38)], applied to finding the ground state of the antiferromagnetic Heisenberg model of Eq. (47).

FIG. 9. Energy as a function of iteration steps for an optimization of the Heisenberg-model ground state using a 10×10 , $D = 4$ PEPS. We compare the energy convergence for different sampling distributions, i.e., MCMC, direct sampling + Markov chain, and direct sampling with different χ_m . The inset shows a zoom of the last 200 iteration steps, highlighting the variability of the energy between subsequent iteration steps and the final converged energies. All runs use $N = 1000$ samples to estimate expectation values and the MCMC algorithm uses $\tau_i = 100$ equilibration sweeps and one sweep between successive samples.

where S_i^x is the spin operator in the x -direction applied on site i and the sum runs over all pairs of nearest neighbors. All algorithms optimize the energy of a $D = 4$ PEPS on a 10×10 square lattice, starting from a random initial PEPS with uniform positive parameters. For the stochastic reconfiguration, we use a diagonal shift $\delta = 0.01$. The stochastic reconfiguration method is clearly advantageous to gradient descent or ADAM, both in the number of iteration steps required to reach a certain accuracy and in the final accuracy the optimization can reach.

2. Energy convergence

We have shown the energy convergence as a function of iteration steps for stochastic gradient descent, ADAM, and stochastic reconfiguration, and found that the latter converges much more rapidly. A natural question to ask is whether the mismatch between p_c and p_ψ also leads to differences in convergence. In Fig. 9, we show the energy convergence of a 10×10 , $D = 4$ PEPS using SR. For the sampling, we use MCMC with $\tau_s = 1$ and $\tau_i = 100$ sweeps, direct sampling combined with a Markov chain with $\tau_s = 1$ and $\tau_i = 100$ samples, direct sampling with $\chi_m = 0$ and direct sampling with $\chi_m = 4$. The first two yield samples according to p_ψ and the other two according to p_c . From this figure, we see that sampling approaches with $\chi_m = 0$ converge slower than MCMC and direct sampling with $\chi_m = 4$. Direct sampling with $\chi_m = 4$, however, leads to the most stable convergence, and the energies reached are comparable with those obtained using MCMC. The convergence of MCMC suffers from unphysical peaks in energy, that can be traced back to poor mixing of the Markov chain sampler. This problem is absent in the direct sampling approaches, as subsequent samples are uncorrelated from each other. This leads to better statistics,

FIG. 10. Relative energy error as a function of wall time spent in the sampling routine for an optimization of the Heisenberg-model ground state using a 10×10 , $D = 4$ PEPS. Both optimizations were performed with the same modest computational resources, i.e., parallelization over eight CPU cores. The MCMC algorithm consists of $\tau_i = 100$ equilibration sweeps and one sweep between successive samples.

also explaining the more stable convergence that can be noted in the inset of Fig. 9.

In Fig. 10, we show the relative energy error of the variational PEPS as a function of the time spent in the sampling routine. We see that the convergence expressed in wall time of the direct sampling procedure is faster than that of MCMC. Also in this figure, the effect of poor sampling sets in the MCMC procedure is visible.

3. Comparison with other optimization algorithms

To show the accuracy of direct sampling combined with gradient updates to optimize PEPS wave functions, we first compare our method with the full update (FU) method, which approximates imaginary-time evolution by Trotter-Suzuki gates and local truncations [7]. We consider a 11×11 PEPS with bond dimensions $D = 2$ and $D = 4$ and variationally find the ground state of the transverse field Ising model, with Hamiltonian

$$\hat{H} = - \sum_{\langle i,j \rangle} \hat{S}_i^z \hat{S}_j^z - g \sum_i \hat{S}_i^x, \quad (49)$$

where g is a coupling constant quantifying the strength of the transverse field. The model has a critical point between a ferromagnetic ordered phase and a disordered phase in the vicinity of $g \approx 3$. In Fig. 11, we show the difference between energies obtained with the direct sampling approach and energies obtained with full update for different values of g . We see that for small bond dimension $D = 2$, the direct sampling approach performs better than full update, especially in the vicinity of the critical point. Second, we compare the accuracy of the ground state energy of the AFH obtained with direct sampling, with those obtained with simple update [6], full update [7], and MCMC-based sampling [11]. In Fig. 12, we show how energies obtained with different optimization algorithms converge as a function of the PEPS bond dimension. For the 4×4 system, we used $\chi_m = 0$ in our sampling algorithm, and for the 10×10 system we used $\chi_m = D$. Using our

FIG. 11. Comparison of energies between direct sampling and full update for a 11×11 PEPS optimized on the TFI in the vicinity of the critical point. We added a linear term in g to the energy to make differences more visible. For $D = 2$, direct sampling performs better compared to FU, while both methods reach similar precision for $D = 4$. For the direct sampling results, $N = 1000$ samples were used together with the SR update method.

sampling approach, we obtain accuracies that are better than SU and FU. Note however that the VMC results were obtained using the $U(1)$ symmetry, which is not the case for the SU and FU results. The direct sampling results are consistent with

FIG. 12. Relative energy error $(E - E_{\text{exact}})/E_{\text{exact}}$ as a function of bond dimension on the ground state of the AFH model for different optimization algorithms and different lattice sizes. An estimation of the exact energy is obtained through QMC. For the direct sampling results, $N = 1000$ samples were used together with the SR update method. Note that direct sampling and MCMC use the $U(1)$ projector to obtain samples with specific $U(1)$ -charge, which may explain the difference in accuracy compared to simple and full update.

MCMC sampling. This shows that, although the sampling distribution p_c is not equal to p_ψ and importance sampling needs to be used, the accuracy of the optimized wave functions does not suffer from the difference in sampling distribution. By employing the environment of the marginal distributions, the distribution p_c can be brought sufficiently close to p_ψ so that variances of sampled quantities are not affected. This was already shown in Fig. 5. It has to be noted that in Ref. [11], it was reported that the sample set size required to obtain reliable gradients was 50–500 times larger than the sample set sizes used in this work. This may be explained by the uncorrelated nature of the samples drawn using direct sampling, compared to the existence of autocorrelation in Markov chain-based sampling approaches. Another important difference with Ref. [11] is that in our simulations the PEPS can be initialized randomly, compared with starting from a PEPS that is already optimized using the SU algorithm. The discrepancy in required sample set size may thus not be related to the direct sampling approach, but rather an indication of the importance of initialization in sampling-based tensor-network algorithms.

VI. CONCLUSION

We have presented an algorithm for drawing uncorrelated samples from a probability distribution described by a double-layer PEPS tensor network. This direct sampling algorithm was inspired by quantum jump Monte Carlo and can be performed in the single layer picture, thus scaling favorably with the bond dimension D of the PEPS. The algorithm can be viewed as sampling a chain of conditional probability distributions, similar to the perfect sampling algorithm for MPS. The absence of a canonical form for a generic PEPS tensor network, however, implies that the marginal probability distributions cannot be constructed trivially. This forces us to employ importance sampling techniques as the sampled distribution p_c is not equal to that defined by the square amplitudes of the PEPS, p_ψ . However, we provide a way to systematically reduce this discrepancy and show that the resulting algorithm provides meaningful samples.

Specifically, we showed by dimensionality reduction that the density of samples drawn according to our algorithm largely overlaps with the density of samples drawn using algorithms sampling p_ψ . We showed that using our samples as transitions in a Markov chain results in small autocorrelation times and large acceptance rates. Our sampling algorithm can be used in a VMC algorithm to determine ground states of quantum many-body systems. The algorithm allows us to obtain expectation values efficiently and reliably, as is reflected in the accuracies we obtain on the transverse field Ising model and the antiferromagnetic Heisenberg model. The resulting optimization is more stable than that using MCMC-based sampling when sample sets of comparable sizes are used.

We believe that our algorithm is useful for future studies using VMC in combination with PEPS wave functions. In particular, we expect that direct sampling will be more efficient than the use of MCMC in large-scale simulations, since the former can be parallelized indefinitely. We can extend our framework beyond ground-state simulations and consider excited states and time evolution; for the latter, we can use a

stochastic implementation of the time-dependent variational principle.

One interesting open question concerns the sign structure of the variational PEPS wave function. In the case of neural-network inspired wave functions, it has been observed that it is crucial to impose the correct sign structure on the wave function for the VMC optimization to be well-behaved [68–70]. In our simulations of the nearest-neighbour Heisenberg model, we have observed that the optimization or variational result does not change when imposing a sign structure or not. It would be very interesting to apply our methods to frustrated systems (as in Ref. [34]), for which the sign structure is not known *a priori*, and investigate this in detail.

Other applications include generative machine learning with PEPS tensor networks, which requires the ability to sample directly from a probability distribution defined by a single layer PEPS [71,72].

A final interesting avenue is the combination of direct sampling with the class of isometric PEPS [13]. The particular structure of isometric PEPS renders the environment of the marginal distributions— $E_m(\mathcal{S}_{>y})$ in Eq. (31)—trivial, and would allow us to directly sample from the distribution p_ψ . This can be combined with VMC and variational optimization for unitary tensor networks manifolds [73] to find ground states that can be represented by isometric PEPS. We will investigate this avenue in future work.

Code to reproduce the results of this paper can be found in Ref. [74].

ACKNOWLEDGMENTS

The authors would like to thank Sheng-Hsuan Lin and Frank Pollmann for inspiring discussions and valuable comments on our manuscript. This work has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme [grant agreements No 715861 (ERQUAF) and 647905 (QUTE)] and from Research Foundation Flanders (FWO) via grants GOE1520N and FWO18/ASP/279. Computational resources (Stevin Supercomputer Infrastructure) and services used in this work were provided by the VSC (Flemish Supercomputer Center), and the Flemish Government: Department EWI.

APPENDIX A: QUANTUM JUMP MONTE CARLO

Quantum jump Monte Carlo originated as a sampling approach for density matrices of systems connected to an environment subject to time evolution according to the Lindblad equation:

$$\frac{d\rho}{dt} = -\frac{i}{\hbar}[H, \rho] - \frac{1}{2} \sum_m (c_m^\dagger c_m \rho + \rho c_m^\dagger c_m - 2c_m \rho c_m^\dagger). \quad (\text{A1})$$

Here, ρ is the density matrix of the system under study, H is the Hamiltonian coupling the degrees of freedom within the system, and c_m are Lindblad operators which describe the interaction of the system with the environment. It is convenient to introduce the effective Hamiltonian $H_{\text{eff}} = H -$

$\frac{i}{2} \sum_m c_m^\dagger c_m$, leading to the more compact equation

$$\frac{d\rho}{dt} = -i(H_{\text{eff}}\rho - \rho H_{\text{eff}}^\dagger) + \sum_m c_m \rho c_m^\dagger. \quad (\text{A2})$$

At first order, the evolution of $\rho(t)$ can be written as

$$\rho(t + dt) = \rho(t) - idt(H_{\text{eff}}\rho - \rho H_{\text{eff}}^\dagger) + dt \sum_m c_m \rho c_m^\dagger. \quad (\text{A3})$$

The density operator $\rho(t)$ can be evolved stochastically using the following algorithm.

(1) At time $t = 0$, sample a pure state $|\phi(t = 0)\rangle$ from the mixed density matrix $\rho(t = 0) = \sum_{\{\phi\}} p_\phi \rho_\phi(0)$, where $\rho_\phi(0) = |\phi(0)\rangle\langle\phi(0)|$ and the sum runs over all pure states in the mixture.

(2) For a small increment dt , evolve $|\phi(t)\rangle$ using one of the following two ways, where $dp = \sum_m \langle\phi(t)|c_m^\dagger c_m|\phi(t)\rangle = \sum_m dp_m$:

(a) with probability $1 - dp$, evolve using

$$|\phi(t + dt)\rangle = \frac{1 - iH_{\text{eff}}dt}{\sqrt{1 - dp}} |\phi(t)\rangle; \quad (\text{A4})$$

(b) with probability dp , evolve using

$$|\phi(t + dt)\rangle = \frac{c_m}{\sqrt{dp_m/dt}} |\phi(t)\rangle, \quad (\text{A5})$$

where a particular c_m is chosen according to the probability distribution dp_m/dp .

(3) Repeat step two until the desired time t is reached.

On the level of the density operator $\rho_\phi(t + dt)$, resulting from evolving the pure state $|\phi(t)\rangle$, the evolution procedure leads to

$$\begin{aligned} \rho_\phi(t + dt) &= |\phi(t)\rangle\langle\phi(t)| - idt(H_{\text{eff}}|\phi(t)\rangle\langle\phi(t)| \\ &\quad - |\phi(t)\rangle\langle\phi(t)|H_{\text{eff}}^\dagger) + dt \sum_m c_m |\phi(t)\rangle\langle\phi(t)|c_m^\dagger. \end{aligned}$$

The sampling over the initial mixture of pure states leads then to the properly evolved density matrix $\rho(t) = \sum_{\{\phi\}} p_\phi \rho_\phi(t)$.

APPENDIX B: SAMPLING OF DENSITY MATRIX

In the sampling algorithm outlined in the main text, we need to sample from the distribution

$$p_c(\mathcal{S}_y | \mathcal{S}_{<y}) = \langle\phi_{y-1}|_b c(\mathcal{S}_y)^\dagger c(\mathcal{S}_y) |\phi_{y-1}\rangle_b, \quad (\text{B1})$$

in the case that $\chi_m = 0$, or

$$p_c(\mathcal{S}_y | \mathcal{S}_{<y}) = \langle\phi_{y-1}|_b c(\mathcal{S}_y)^\dagger E_m(\mathcal{S}_{>y}) c(\mathcal{S}_y) |\phi_{y-1}\rangle_b, \quad (\text{B2})$$

in the case that the environment of the marginal is used. For notational simplicity, we will use the latter notation, assuming the MPO $E_m(\mathcal{S}_{>y})$ is the unit MPO when $\chi_m = 0$ is used. In diagrammatic notation, this probability can be drawn as

$$p_c(\mathcal{S}_y | \mathcal{S}_{<y}) = \begin{array}{c} \text{Diagram: A 3x4 grid of tensors. The top and bottom rows consist of orange squares, and the middle row consists of green squares. Horizontal and vertical lines connect adjacent squares. The four orange squares in the top row are connected to four circles labeled } s_1, s_2, s_3, s_4 \text{ from left to right. Similarly, the four orange squares in the bottom row are connected to four circles labeled } s_1, s_2, s_3, s_4 \text{ from left to right.} \end{array}, \quad (\text{B3})$$

where we define $\mathcal{S}_y \equiv (s_1, \dots, s_N)$ for notational simplicity.

A multidimensional probability distribution $p(s_1, s_2, s_3, s_4)$ such as the one in Eq. (B3) can always be decomposed in a product of conditional probabilities using the bayesian rule:

$$p(s_1, s_2, s_3, s_4) = p(s_4 | s_3, s_2, s_1) p(s_3 | s_2, s_1) p(s_2 | s_1) p(s_1), \quad (\text{B4})$$

where the marginal probability distribution is given by

$$p(s_1) = \sum_{s_1, s_2, s_3} p(s_1, s_2, s_3, s_4), \quad (\text{B5})$$

and $p(x|y)$ is the conditional probability distribution of the set of degrees of freedom x , conditioned on the set of degrees of freedom y . Sampling from $p(s_1, s_2, s_3, s_4)$ can now be done sequentially, by first sampling a configuration (s_1) from $p(s_1)$, then using the outcome to construct $p(s_2 | s_1)$ and sample from this the configuration (s_1, s_2) , and repeating this until a configuration of all degrees of freedom is sampled.

Any probability distribution can be written in this form, but calculating the marginal probability distributions is in principle exponentially hard. To sample this probability distribution, we adapt the perfect sampling approach for unitary tensor networks [45]. To this end, we first calculate $p(s_1)$ defined in Eq. (B4) as the marginal over all trailing degrees of freedom except the first:

$$p_c(s_1 | \mathcal{S}_{<y}) = \sum_{\{s_2 \dots s_4\}} p_c(\mathcal{S}_y | \mathcal{S}_{<y})$$

$$= \begin{array}{c} \text{Diagram: A vertical stack of three tensors. The top and bottom are orange squares, and the middle is a green square. They are connected vertically. The top orange square is connected to a circle labeled } s_1. \text{ The bottom orange square is also connected to a circle labeled } s_1. \text{ The middle green square is connected to a circle labeled } e_1. \end{array}. \quad (\text{B6})$$

Here, we defined the environment e_1 obtained by contracting all columns of the tensor network defined in Eq. (B3), where we also contract over the physical degrees of freedom to perform the marginalization. The contraction in Eq. (B6) yields (the diagonal of) a density matrix, from which we sample s_1 according to the probabilities defined as its entries. As this is a low-dimensional distribution, we can make use of simple algorithms yielding uncorrelated samples. Once we have sampled s_1 , we can construct the probability distribution $p(s_2 | s_1)$ as follows

$$p_c(s_2 | \mathcal{S}_{\leq 1, < y}) = \sum_{s_3, s_4} p_c(s_2, s_3, s_4 | \mathcal{S}_{\leq 1, < y})$$

$$= \begin{array}{c} \text{Diagram: A 3x2 grid of tensors. The top and bottom rows consist of orange squares, and the middle row consists of green squares. Horizontal and vertical lines connect adjacent squares. The two orange squares in the top row are connected to two circles labeled } s_1, s_2 \text{ from left to right. Similarly, the two orange squares in the bottom row are connected to two circles labeled } s_1, s_2 \text{ from left to right. The middle green square is connected to a circle labeled } e_2. \end{array}$$

$$= \begin{array}{c} \text{Diagram: A circular contraction. A central green square is connected to four orange squares arranged in a circle. The top orange square is connected to a circle labeled } s_2. \text{ The bottom orange square is also connected to a circle labeled } s_2. \text{ The left orange square is connected to a circle labeled } \sigma_1. \text{ The right orange square is connected to a circle labeled } e_2. \end{array}. \quad (\text{B7})$$

Here, we require the environment e_2 , which is defined in a similar manner as e_1 , but now contracting only over all but the first two columns in Eq. (B3). Note that during calculation of e_1 , we already calculated all other environments $e_{2 \dots 4}$. Furthermore, we compute the environment σ_1 by contracting the first column of Eq. (B3), where the orange tensors are projected on the value of s_1 sampled in the previous step. The resulting contraction yields a density matrix from which we can sample s_2 using the probabilities specified on the diagonal. We can keep repeating this procedure until we have sampled all the degrees of freedom. In total, this yields an uncorrelated configuration from the probability distribution $p_c(\mathcal{S}_y)$.

- [1] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, Cambridge, 2017).
- [2] R. G. Melko, G. Carleo, J. Carrasquilla, and J. I. Cirac, Restricted boltzmann machines in quantum physics, *Nat. Phys.* **15**, 887 (2019).
- [3] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011), january 2011 Special Issue.
- [4] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions, [arXiv:cond-mat/0407066](https://arxiv.org/abs/cond-mat/0407066).
- [5] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, Classical Simulation of Infinite-Size Quantum Lattice Systems in Two Spatial Dimensions, *Phys. Rev. Lett.* **101**, 250602 (2008).
- [6] H. C. Jiang, Z. Y. Weng, and T. Xiang, Accurate Determination of Tensor Network State of Quantum Lattice Models in Two Dimensions, *Phys. Rev. Lett.* **101**, 090603 (2008).
- [7] M. Lubasch, J. I. Cirac, and M. C. Bañuls, Algorithms for finite projected entangled pair states, *Phys. Rev. B* **90**, 064425 (2014).
- [8] H. N. Phien, J. A. Bengua, H. D. Tuan, P. Corboz, and R. Orús, Infinite projected entangled pair states algorithm improved: Fast full update and gauge fixing, *Phys. Rev. B* **92**, 035142 (2015).
- [9] L. Vanderstraeten, J. Haegeman, P. Corboz, and F. Verstraete, Gradient methods for variational optimization of projected entangled-pair states, *Phys. Rev. B* **94**, 155123 (2016).
- [10] P. Corboz, Variational optimization with infinite projected entangled-pair states, *Phys. Rev. B* **94**, 035133 (2016).
- [11] W. Y. Liu, S. J. Dong, Y. J. Han, G. C. Guo, and L. He, Gradient optimization of finite projected entangled pair states, *Phys. Rev. B* **95**, 195154 (2017).
- [12] H.-J. Liao, J.-G. Liu, L. Wang, and T. Xiang, Differentiable Programming Tensor Networks, *Phys. Rev. X* **9**, 031041 (2019).
- [13] M. P. Zaletel and F. Pollmann, Isometric Tensor Network States in Two Dimensions, *Phys. Rev. Lett.* **124**, 037201 (2020).
- [14] P. Corboz and F. Mila, Tensor network study of the shastry-sutherland model in zero magnetic field, *Phys. Rev. B* **87**, 115144 (2013).
- [15] P. Corboz, T. M. Rice, and M. Troyer, Competing States in the t - j Model: Uniform d -Wave State Versus Stripe State, *Phys. Rev. Lett.* **113**, 046402 (2014).
- [16] B.-X. Zheng, C.-M. Chung, P. Corboz, G. Ehlers, M.-P. Qin, R. M. Noack, H. Shi, S. R. White, S. Zhang, and G. K.-L. Chan, Stripe order in the underdoped region of the two-dimensional hubbard model, *Science* **358**, 1155 (2017).
- [17] J.-Y. Chen, L. Vanderstraeten, S. Capponi, and D. Poilblanc, Non-Abelian chiral spin liquid in a quantum antiferromagnet revealed by an ipeps study, *Phys. Rev. B* **98**, 184409 (2018).
- [18] P. Czarnik, L. Cincio, and J. Dziarmaga, Projected entangled pair states at finite temperature: Imaginary time evolution with ancillas, *Phys. Rev. B* **86**, 245101 (2012).
- [19] P. Czarnik and J. Dziarmaga, Variational approach to projected entangled pair states at finite temperature, *Phys. Rev. B* **92**, 035152 (2015).
- [20] A. Kshetrimayum, H. Weimer, and R. Orús, A simple tensor network algorithm for two-dimensional steady states, *Nat. Commun.* **8**, 1291 (2017).
- [21] L. Vanderstraeten, M. Mariën, F. Verstraete, and J. Haegeman, Excitations and the tangent space of projected entangled-pair states, *Phys. Rev. B* **92**, 201111(R) (2015).
- [22] L. Vanderstraeten, J. Haegeman, and F. Verstraete, Simulating excitation spectra with projected entangled-pair states, *Phys. Rev. B* **99**, 165121 (2019).
- [23] B. Ponsioen and P. Corboz, Excitations with projected entangled pair states using the corner transfer matrix method, *Phys. Rev. B* **101**, 195109 (2020).
- [24] P. Czarnik and J. Dziarmaga, Time evolution of an infinite projected entangled pair state: An algorithm from first principles, *Phys. Rev. B* **98**, 045110 (2018).
- [25] P. Czarnik, J. Dziarmaga, and P. Corboz, Time evolution of an infinite projected entangled pair state: An efficient algorithm, *Phys. Rev. B* **99**, 035115 (2019).
- [26] C. Hubig and J. I. Cirac, Time-dependent study of disordered models with infinite projected entangled pair states, *Sci. Post. Phys.* **6**, 31 (2019).
- [27] A. W. Sandvik and G. Vidal, Variational Quantum Monte Carlo Simulations with Tensor-Network States, *Phys. Rev. Lett.* **99**, 220602 (2007).
- [28] A. W. Sandvik, Scale-Renormalized Matrix-Product States for Correlated Quantum Systems, *Phys. Rev. Lett.* **101**, 140603 (2008).
- [29] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Simulation of Quantum Many-Body Systems with Strings of Operators and Monte Carlo Tensor Contractions, *Phys. Rev. Lett.* **100**, 040501 (2008).
- [30] L. Wang, I. Pizorn, and F. Verstraete, Monte Carlo simulation with Tensor Network States, *Phys. Rev. B* **83**, 134421 (2011).
- [31] C.-P. Chou, F. Pollmann, and T.-K. Lee, Matrix-product-based projected wave functions ansatz for quantum many-body ground states, *Phys. Rev. B* **86**, 041105(R) (2012).
- [32] O. Sikora, H.-W. Chang, C.-P. Chou, F. Pollmann, and Y.-J. Kao, Variational monte carlo simulations using tensor-product projected states, *Phys. Rev. B* **91**, 165113 (2015).
- [33] W.-Y. Liu, Y.-Z. Huang, S.-S. Gong, and Z.-C. Gu, Accurate simulation for finite projected entangled pair states in two dimensions, *Phys. Rev. B* **103**, 235155 (2021).
- [34] W.-Y. Liu, S.-S. Gong, Y.-B. Li, D. Poilblanc, W.-Q. Chen, and Z.-C. Gu, Gapless quantum spin liquid and global phase diagram of the spin-1/2 J_1 - J_2 square antiferromagnetic Heisenberg model, [arXiv:2009.01821](https://arxiv.org/abs/2009.01821).
- [35] G. E. P. Box and M. E. Muller, A Note on the Generation of Random Normal Deviates, *Ann. Math. Stat.* **29**, 610 (1958).
- [36] L. Dinh, J. Sohl-Dickstein, and S. Bengio, Density estimation using Real NVP, in *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings* (ICLR, 2016).
- [37] D. P. Kingma and P. Dhariwal, Glow: Generative Flow with Invertible 1x1 Convolutions, *Advances in Neural Information Processing Systems* **2018-December**, 10215 (2018).
- [38] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu, Pixel recurrent neural networks, in *Proceedings of The 33rd International Conference on Machine Learning*, Vol. 48, edited by M. F. Balcan and K. Q. Weinberger (PMLR, New York, 2016), pp. 1747–1756.

- [39] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, WaveNet: A generative model for raw audio, [arXiv:1609.03499](#).
- [40] D. Wu, L. Wang, and P. Zhang, Solving Statistical Mechanics Using Variational Autoregressive Networks, *Phys. Rev. Lett.* **122**, 080602 (2019).
- [41] M. S. Albergo, G. Kanwar, and P. E. Shanahan, Flow-based generative models for Markov chain Monte Carlo in lattice field theory, *Phys. Rev. D* **100**, 034515 (2019).
- [42] O. Sharir, Y. Levine, N. Wies, G. Carleo, and A. Shashua, Deep Autoregressive Models for the Efficient Variational Simulation of Many-Body Quantum Systems, *Phys. Rev. Lett.* **124**, 020503 (2020).
- [43] M. Hibat-Allah, M. Ganahl, L. E. Hayward, R. G. Melko, and J. Carrasquilla, Recurrent neural network wave functions, *Phys. Rev. Research* **2**, 023358 (2020).
- [44] C. Casert, T. Vieijra, S. Whitelam, and I. Tamblyn, Dynamical Large Deviations of Two-Dimensional Kinetically Constrained Models Using a Neural-Network State Ansatz, *Phys. Rev. Lett.* **127**, 120602 (2021).
- [45] A. J. Ferris and G. Vidal, Perfect sampling with unitary tensor networks, *Phys. Rev. B* **85**, 165146 (2012).
- [46] M. Frías-Pérez, M. Mariën, D. P. García, M. C. Bañuls, and S. Iblisdir, Collective Monte Carlo updates through tensor network renormalization, [arXiv:2104.13264](#).
- [47] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Computational Complexity of Projected Entangled Pair States, *Phys. Rev. Lett.* **98**, 140506 (2007).
- [48] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Unifying projected entangled pair state contractions, *New J. Phys.* **16**, 033014 (2014).
- [49] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.* **21**, 1087 (1953).
- [50] W. K. Hastings, Monte carlo sampling methods using Markov chains and their applications, *Biometrika* **57**, 97 (1970).
- [51] U. Wolff, Collective Monte Carlo Updating for Spin Systems, *Phys. Rev. Lett.* **62**, 361 (1989).
- [52] A. J. Daley, Quantum trajectories and open many-body quantum systems, *Adv. Phys.* **63**, 77 (2014).
- [53] D. P. Kingma and J. Ba, ADAM: A method for stochastic optimization, [arXiv:1412.6980](#).
- [54] S. Sorella, Green Function Monte Carlo with Stochastic Reconfiguration, *Phys. Rev. Lett.* **80**, 4558 (1998).
- [55] E. Neuscamman, C. J. Umrigar, and G. K. L. Chan, Optimizing large parameter sets in variational quantum Monte Carlo, *Phys. Rev. B* **85**, 045103 (2012).
- [56] S. I. Amari, Natural gradient works efficiently in learning, *Neural Comput.* **10**, 251 (1998).
- [57] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Unifying time evolution and optimization with matrix product states, *Phys. Rev. B* **94**, 165116 (2016).
- [58] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Time-Dependent Variational Principle for Quantum Lattices, *Phys. Rev. Lett.* **107**, 070601 (2011).
- [59] L. Vanderstraeten, J. Haegeman, and F. Verstraete, Tangent-space methods for uniform matrix product states, *Sci. Post Phys. Lect. Notes*, 7 (2019).
- [60] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Criticality, the Area Law, and the Computational Power of Projected Entangled Pair States, *Phys. Rev. Lett.* **96**, 220601 (2006).
- [61] K. Pearson, LIII. On lines and planes of closest fit to systems of points in space, *London, Edinburgh, Dublin Philosophical Magazine J. Sci.* **2**, 559 (1901).
- [62] K. Binder and A. P. Young, Spin glasses: Experimental facts, theoretical concepts, and open questions, *Rev. Mod. Phys.* **58**, 801 (1986).
- [63] R. H. Swendsen and J.-S. Wang, Replica Monte Carlo Simulation of Spin-Glasses, *Phys. Rev. Lett.* **57**, 2607 (1986).
- [64] S. Jensen, N. Read, and A. P. Young, Nontrivial maturation metastate-average state in a one-dimensional long-range ising spin glass: Above and below the upper critical range, *Phys. Rev. E* **104**, 034105 (2021).
- [65] C. Gattringer and C. B. Lang, *Quantum Chromodynamics on the Lattice* (Springer, Berlin, 2010), Vol. 788.
- [66] A. Kong, A note on importance sampling using standardized weights, University of Chicago, Dept. of Statistics, Tech. Rep **348** (1992).
- [67] J. S. Liu, Metropolized independent sampling with comparisons to rejection sampling and importance sampling, *Statistics computing* **6**, 113 (1996).
- [68] K. Choo, T. Neupert, and G. Carleo, Two-dimensional frustrated J_1 - J_2 model studied with neural network quantum states, *Phys. Rev. B* **100**, 125124 (2019).
- [69] T. Westerhout, N. Astrakhantsev, K. S. Tikhonov, M. I. Katsnelson, and A. A. Bagrov, Generalization properties of neural network approximations to frustrated magnet ground states, *Nat. Commun.* **11**, 1593 (2020).
- [70] A. Szabó and C. Castelnovo, Neural network wave functions and the sign problem, *Phys. Rev. Research* **2**, 033075 (2020).
- [71] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, Unsupervised Generative Modeling using Matrix Product States, *Phys. Rev. X* **8**, 031012 (2018).
- [72] S. Cheng, L. Wang, T. Xiang, and P. Zhang, Tree tensor networks for generative modeling, *Phys. Rev. B* **99**, 155131 (2019).
- [73] M. Hauru, M. Van Damme, and J. Haegeman, Riemannian optimization of isometric tensor networks, *Sci. Post Phys.* **10**, 040 (2021).
- [74] <https://github.com/tvieijra/DirectSamplingPEPS.jl>.