

REVIEW

The density-matrix renormalization group: a short introduction

BY ULRICH SCHOLLWÖCK^{1,2*}

¹*Department of Physics, Arnold Sommerfeld Center for Theoretical Physics
and Center for NanoScience, University of Munich, Theresienstrasse 37,
80333 Munich, Germany*

²*Institute for Advanced Study Berlin, Wallotstrasse 19, 14159 Berlin, Germany*

The density-matrix renormalization group (DMRG) method has established itself over the last decade as the leading method for the simulation of the statics and dynamics of one-dimensional strongly correlated quantum lattice systems. The DMRG is a method that shares features of a renormalization group procedure (which here generates a flow in the space of reduced density operators) and of a variational method that operates on a highly interesting class of quantum states, so-called matrix product states (MPSs). The DMRG method is presented here entirely in the MPS language. While the DMRG generally fails in larger two-dimensional systems, the MPS picture suggests a straightforward generalization to higher dimensions in the framework of tensor network states. The resulting algorithms, however, suffer from difficulties absent in one dimension, apart from a much more unfavourable efficiency, such that their ultimate success remains far from clear at the moment.

Keywords: density-matrix renormalization group; matrix product states; entanglement

1. Introduction

In the last decade, the density-matrix renormalization group (DMRG), invented by White [1,2], has established itself as the leading method for the simulation of the statics and dynamics of one-dimensional strongly correlated quantum systems, both at zero and finite temperatures (for more recent reviews, see [3–5]). Nevertheless, attempts have been made to extend this method to the simulation of two-dimensional quantum systems and (three-dimensional) small clusters, such as nuclei or molecules. These attempts are currently being advanced owing to the alternative view of the DMRG as a variational method in the space of matrix product states (MPSs) [6–9] because they allow for an easy extension of the ansatz class to higher dimensions [10,11], which might hold interesting promise for such finite clusters of particles. In this short overview, following [5], I will first introduce and discuss MPSs, introduce the most basic version of the DMRG as a variational

*schollwoeck@lmu.de

One contribution of 11 to a Theme Issue ‘New applications of the renormalization group in nuclear, particle and condensed matter physics’.

method and discuss its dual nature as a variational and renormalization group (RG) method, as well as the reasons for its successes and limitations. I will conclude by a short discussion of the generalizations to higher dimensions.

2. Density-matrix renormalization group in one dimension

(a) Matrix product states

Let us assume that our system consists of L sites on a one-dimensional chain with local state spaces $\{|\sigma_i\rangle\}$ on site i of dimension d each. Then, it can be shown (e.g. [5]) that any quantum state

$$|\psi\rangle = \sum_{\sigma} c_{\sigma} |\sigma\rangle, \quad (2.1)$$

where $|\sigma\rangle = |\sigma_1, \sigma_2, \dots, \sigma_L\rangle$, can be brought (for example) by a sequence of singular value decompositions (SVDs) exactly into the form of an MPS, namely

$$|\psi\rangle = \sum_{\sigma} M^{\sigma_1} M^{\sigma_2} \dots M^{\sigma_{L-1}} M^{\sigma_L} |\sigma\rangle, \quad (2.2)$$

where on each lattice site i we introduce a set of d matrices M^{σ_i} . To see this, let us briefly recall the key properties of an SVD (e.g. [12]): for an arbitrary, rectangular matrix A of dimension $(m \times n)$, it provides a decomposition $A = USV^{\dagger}$, where the matrices U , S and V^{\dagger} have dimensions $(m \times p)$, $(p \times p)$ and $(p \times n)$, where $p = \min(m, n)$, the smaller of the two original dimensions. In addition, the three matrices have the following properties: U consists of p orthonormal columns, which can be read as a set of orthonormal vectors (hence, $U^{\dagger}U = I$); similarly, V^{\dagger} consists of p orthonormal rows, which again can be read as a set of orthonormal vectors (and $V^{\dagger}V = I$); they are referred to as the left and right singular vectors. S is a non-negative diagonal matrix, with $r < p$ positive singular values on the diagonal; r is the rank of A . In order to carry out the transformation into an MPS, we start by reshaping c_{σ} into a matrix $\Psi_{\sigma_1, (\sigma_2 \dots \sigma_L)} = c_{\sigma_1 \dots \sigma_L}$ of dimension $(d \times d^{L-1})$. This is now singular value decomposed as

$$\Psi_{\sigma_1, (\sigma_2 \dots \sigma_L)} = \sum_{a_1} U_{\sigma_1, a_1} S_{a_1} V_{a_1, (\sigma_2 \dots \sigma_L)}^{\dagger}. \quad (2.3)$$

The U -matrix has dimension $d \times d$. As the next step, we reshape $S_{a_1} V_{a_1, (\sigma_2 \dots \sigma_L)}^{\dagger}$ into a matrix $\Psi_{(a_1 \sigma_2), (\sigma_3 \dots \sigma_L)}$ of dimension $(d^2 \times d^{L-2})$, where the row dimension may be smaller if singular values of the previous SVD happen to be zero. This is again singular value decomposed as

$$\Psi_{(a_1 \sigma_2), (\sigma_3 \dots \sigma_L)} = \sum_{a_2} U_{(a_1 \sigma_2), a_2} S_{a_2} V_{a_2, (\sigma_3 \dots \sigma_L)}^{\dagger}, \quad (2.4)$$

with the dimension of the U -matrix being $d^2 \times d^2$. If we continue this procedure to the end of the chain, we end up with a sequence of U -matrices as U_{σ_1, a_1} ,




Figure 1. Definition of the graphical representation of a matrix product state (MPS) constituent, a matrix M_{ab}^{σ} , its conjugate, and of an MPS. Note that on the first and last sites, the matrices reduce to row and column vectors (one horizontal leg only). The evaluation rule is to contract over all connected legs.

$U_{(a_1\sigma_2),a_2}$, $U_{(a_2\sigma_3),a_3}$ and so on, with the last matrix being $U_{(a_{L-1}\sigma_L),1}$. The total state then reads

$$|\psi\rangle = \sum_{\sigma} \sum_{a_1, \dots, a_{L-1}} U_{\sigma_1, a_1} U_{(a_1\sigma_2), a_2} \dots U_{(a_{L-2}\sigma_{L-1}), a_{L-1}} U_{(a_{L-1}\sigma_L), 1} |\sigma\rangle. \quad (2.5)$$

To simplify the notation, we now define the d matrices M^{σ_i} per site as $M_{a_{i-1}, a_i}^{\sigma_i} = U_{(a_{i-1}\sigma_i), a_i}$ and use that the sum over the a_i is then nothing but a sequence of matrix multiplications of M^{σ_i} matrices as in equation (2.2).

It is useful to introduce a graphical notation for MPSs (figure 1), where each matrix is represented by a blob with outgoing legs, namely a vertical physical leg and horizontal legs corresponding to the matrix indices. The rule is that all connected legs are contracted, leading to the MPSs as defined above.

If we review the above sequence of decompositions and the sizes of the U -matrices, the resulting M^{σ} matrix dimensions will be $(1 \times d), (d \times d^2), \dots, (d^{L/2-1} \times d^{L/2}), (d^{L/2} \times d^{L/2-1}), \dots, (d^2 \times d), (d \times 1)$: the M -matrices are always cut in the row dimension by a factor of d by the reshaping from the U -matrices; beyond the centre of the chain, the column dimension will be smaller than the row dimension of the singular value decomposed matrix, hence the dimensions of the U -matrices will go down again. In systems of interest, where $d^{L/2}$ will be exponentially large in system size, this exact representation is obviously useless in numerical practice. But let us assume that we consider an MPS where matrix dimensions do not exceed some number D , which will be $O(100)$ through $O(1000)$, and that these approximate MPSs are a very close approximation to the exact quantum states of interest. Under this assumption, we may ask about quantum mechanics in this restricted set of MPSs of dimension D .

Before we consider typical MPS operations, we note that MPS representations of quantum states are not unique: between any matrix pair $M^{\sigma_i} M^{\sigma_{i+1}}$, we may insert $I = XX^{-1}$, where X is an invertible matrix of suitable dimensions, changing $M^{\sigma_i} \rightarrow M^{\sigma_i} X$ and $M^{\sigma_{i+1}} \rightarrow X^{-1} M^{\sigma_{i+1}}$. One of the main usages of this is to bring the MPS into *canonical form*: any MPS can be written in left-canonical form

$$|\psi\rangle = \sum_{\sigma} A^{\sigma_1} A^{\sigma_2} \dots A^{\sigma_{L-1}} A^{\sigma_L} |\sigma\rangle, \quad (2.6)$$

where the matrices A^{σ_i} are *left-normalized*, $\sum_{\sigma_i} A^{\sigma_i \dagger} A^{\sigma_i} = I$, or in right-canonical form

$$|\psi\rangle = \sum_{\sigma} B^{\sigma_1} B^{\sigma_2} \dots B^{\sigma_{L-1}} B^{\sigma_L} |\sigma\rangle, \quad (2.7)$$

where the matrices B^{σ_i} are *right-normalized*, $\sum_{\sigma_i} B^{\sigma_i} B^{\sigma_i \dagger} = I$, or in mixed-canonical form

$$|\psi\rangle = \sum_{\sigma} A^{\sigma_1} A^{\sigma_2} \dots A^{\sigma_i} \Psi B^{\sigma_{i+1}} \dots B^{\sigma_{L-1}} B^{\sigma_L} |\sigma\rangle, \tag{2.8}$$

where the matrix Ψ is of matching dimensions; one can also contract e.g. $A^{\sigma_i} \Psi \rightarrow M^{\sigma_i}$, which generates a variety of ansatz classes for the DMRG. To see how this can be done iteratively, consider a state that is already partially in left-canonical form on its left end,

$$|\psi\rangle = \sum_{\sigma} A^{\sigma_1} A^{\sigma_2} \dots A^{\sigma_i} M^{\sigma_{i+1}} \dots M^{\sigma_L} |\sigma\rangle. \tag{2.9}$$

Then, one way of making $M^{\sigma_{i+1}}$ left-canonical is to reshape the set of matrices $M^{\sigma_{i+1}}_{a_i, a_{i+1}}$ into one matrix $M_{(\sigma_{i+1} a_i), a_{i+1}}$, which upon an SVD reads $\sum_{s_{i+1}} U_{(\sigma_{i+1} a_i), s_{i+1}} S_{s_{i+1}} V^{\dagger}_{s_{i+1}, a_{i+1}}$. U is reshaped into a set of matrices $A^{\sigma_i}_{a_i, s_{i+1}}$, which meet the left-normalization condition owing to the column orthonormality of U . SV^{\dagger} is multiplied into $M^{\sigma_{i+2}}$, and the next iterative step occurs at site $i + 2$. Similarly, one can start from the right to obtain right-normalized matrices. In fact, the proof of the representation in equation (2.2) generated left-normalized matrices right away, owing to the orthonormality properties of the U -matrices.

If we form states $|a_i\rangle_A$ on block A formed from sites 1 through to i and $|a_i\rangle_B$ on block B formed from sites $i + 1$ through to L as

$$|a_i\rangle_A = \sum_{\sigma_1, \dots, \sigma_i} (M^{\sigma_1} \dots M^{\sigma_i})_{a_i} |\sigma_1, \dots, \sigma_i\rangle, \tag{2.10}$$

and similarly for $|a_i\rangle_B$, then the left- and right-normalization conditions correspond to orthonormality conditions ${}_A\langle a_i | a'_i \rangle_A = \delta_{a_i, a'_i}$ and ${}_B\langle a_i | a'_i \rangle_B = \delta_{a_i, a'_i}$, respectively. Generally, it is true that one can write

$$|a_{i+1}\rangle_A = \sum_{a_i, \sigma_{i+1}} M^{\sigma_{i+1}}_{a_i, a_{i+1}} |a_i\rangle_A |\sigma_{i+1}\rangle, \tag{2.11}$$

where $M^{\sigma_{i+1}}_{a_i, a_{i+1}} = \langle a_i \sigma_{i+1} | a_{i+1} \rangle$. This establishes the connection between MPS notation and block growth $i \rightarrow i + 1$, which—taking into account decimation of the states $|a_{i+1}\rangle_A$ —makes the connection to RG growth steps.

In the mixed-canonical representation, the nature of the MPS as a systematic low-entanglement approximation to a quantum state becomes most obvious. If we do an SVD on $\Psi = USV^{\dagger}$ and absorb matrices U and V^{\dagger} into the adjoining A - and B -matrices, this does not change the orthonormality of the block states because of the properties of U and V^{\dagger} , but the state now takes the form

$$|\psi\rangle = \sum_{a_i}^r s_{a_i} |a_i\rangle_A |a_i\rangle_B, \tag{2.12}$$

where the sum runs over the r non-vanishing (positive) diagonal elements s_{a_i} of S , linking up states in A and B in pairwise fashion. This is the so-called Schmidt decomposition of a quantum state, with r being the Schmidt rank. If we have to cut the sum at some $D < r$, the quality of the approximation will depend on how




Figure 2. Graphical representation of the calculation of the norm of a state, potentially including an operator whose expectation value is calculated. Lines are contracted over iteratively, starting either from the left or right, for maximal efficiency.

rapidly the s_{a_i} decay; the optimal approximation is given by keeping the D largest s_{a_i} . But it is easy to see that at the same time, the reduced density operators for A and B are given as $\hat{\rho}_A = \text{Tr}_B |\psi\rangle\langle\psi| = \sum_{a_i} s_{a_i}^2 |a_i\rangle_A \langle a_i|$ and analogously $\hat{\rho}_B$. A good approximation is then possible if the eigenspectrum of the reduced density operators decays quickly. It is usually not known, but we may argue that a quickly decaying spectrum corresponds to small entanglement $S_{A|B} = -\sum_a s_a^2 \log_2 s_a^2$, and MPSs systematically approximate low-entanglement states well. Let us take that for granted and postpone the crucial question when this is actually the case to later, and focus on operations on MPSs.

The most important operations on MPSs are overlaps and expectation values. Let us consider the calculation of $\langle\phi|\psi\rangle$, the generalization to $\langle\psi|\hat{O}|\psi\rangle$ is completely obvious. We have

$$\langle\phi|\psi\rangle = \sum_{\sigma} \tilde{M}^{\sigma_1*} \dots \tilde{M}^{\sigma_L*} M^{\sigma_1} \dots M^{\sigma_L}. \tag{2.13}$$

Transposing the scalar formed from the $\tilde{M} \dots \tilde{M}$ (which is the identity operation), we arrive at adjoints with reversed ordering,

$$\langle\phi|\psi\rangle = \sum_{\sigma} \tilde{M}^{\sigma_L\dagger} \dots \tilde{M}^{\sigma_1\dagger} M^{\sigma_1} \dots M^{\sigma_L}. \tag{2.14}$$

To represent this graphically, we introduce an inverted matrix-blob to correspond to the complex-conjugate matrix M^{σ_i*} (figure 1). In the resulting pictorial representation of the overlap calculation (figure 2), this calculation then becomes much simpler, if we follow the rule that all bond indices are summed over. To evaluate this expression efficiently, the following order of contractions is advisable:

$$\langle\phi|\psi\rangle = \sum_{\sigma_L} \tilde{M}^{\sigma_L\dagger} \left(\dots \left(\sum_{\sigma_2} \tilde{M}^{\sigma_2\dagger} \left(\sum_{\sigma_1} \tilde{M}^{\sigma_1\dagger} M^{\sigma_1} \right) M^{\sigma_2} \right) \dots \right) M^{\sigma_L}, \tag{2.15}$$

which contracts in $O(LD^3)$ operations. In order to calculate an expectation value, say at site i , the sum to be evaluated there changes to

$$\sum_{\sigma_i, \sigma'_i} \hat{O}^{\sigma_i, \sigma'_i} M^{\sigma_i*} E M^{\sigma'_i}, \tag{2.16}$$

where E is the result of all the contractions up to site $i - 1$.

If we exploit left-normalization, it is easy to see that for $\langle \psi | \psi \rangle$ matrix contractions collapse trivially to the identity I (similarly for right-normalization, if we rearrange the contractions to go from L to 1). In mixed-canonical notation, expectation values for operators that sit right next to the point where left- and right-normalized matrices meet become particularly efficient, as essentially all contractions become trivial.

The last important operation on an MPS is its *compression*. Assume that some operation (like the application of a matrix product operator (MPO) to an MPS, see below) generates an MPS of dimension D' , but we can only handle dimension D efficiently. How do we approximate $|\psi\rangle$ by $|\tilde{\psi}\rangle$, an MPS of dimension D , such that the distance between the two states becomes minimal, i.e. the approximation becomes optimal? Essentially two techniques exist, one based on SVD and not necessarily optimal, but frequently very close to it, and a variational technique, that is more complicated, but optimal.

Let us consider an MPS in mixed-canonical representation, with $A^{[\ell]} = \text{diag}(s_1, s_2, \dots)$,

$$|\psi\rangle = \sum_{\sigma} A^{\sigma_1} A^{\sigma_2} \dots A^{\sigma_{\ell}} A^{[\ell]} B^{\sigma_{\ell+1}} \dots B^{\sigma_{L-1}} B^{\sigma_L} |\sigma\rangle, \tag{2.17}$$

from which we read off the Schmidt decomposition $|\psi\rangle = \sum_{a_{\ell}=1}^{D'} s_{a_{\ell}} |a_{\ell}\rangle_A |a_{\ell}\rangle_B$, where the states on A and B form orthonormal sets, respectively. We now look for the state $|\tilde{\psi}\rangle$ that approximates $|\psi\rangle$ best in the 2-norm and can be spanned by a smaller number of D states each in A and B . This is achieved by retaining the D largest $s_{a_{\ell}}$, and the compressed state simply reads $|\psi\rangle = \sum_{a_{\ell}=1}^D s_{a_{\ell}} |a_{\ell}\rangle_A |a_{\ell}\rangle_B$. If normalization is desired, the retained singular values must be rescaled.

This procedure rests on the orthonormality of the states on A and B , therefore can only be carried out at one bond. In order to shrink the state on all sites, we have to work our way through all mixed-canonical representations, say from right to left, truncate and shift the boundary between left- and right-normalized matrices by one site to the left, using techniques from canonization.

After the first step of right-canonization of a left-canonical state, it reads

$$|\psi^{(L-1)}\rangle = \sum_{\sigma} A^{\sigma_1} \dots A^{\sigma_{L-1}} U S B^{\sigma_L} |\sigma\rangle, \tag{2.18}$$

where I have already reshaped B , which is right-normalized and guarantees that states formed as $|a_{L-1}\rangle_B = \sum_{\sigma_L} (B^{\sigma_L})_{a_{L-1},1} |\sigma_L\rangle$ are orthonormal. But so are the states

$$|a_{L-1}\rangle_A = \sum_{\sigma_1 \dots \sigma_{L-1}} (A^{\sigma_1} \dots A^{\sigma_{L-1}} U)_{1,a_{L-1}} |\sigma_1 \dots \sigma_{L-1}\rangle, \tag{2.19}$$

as SVD guarantees $U^\dagger U = 1$: we are just doing a basis transformation within the orthonormal basis set constructed from the left-normalized A^{σ_i} . Hence, we have a correct Schmidt decomposition as

$$|\psi^{(L-1)}\rangle = \sum_{a_{L-1}} s_{a_{L-1}} |a_{L-1}\rangle_A |a_{L-1}\rangle_B. \tag{2.20}$$

The difference to a right canonization is now the truncation: matrices U , S and B^{σ_L} are truncated (and singular values possibly renormalized) to \tilde{U} , \tilde{S} and \tilde{B}^{σ_L} just as explained before: retain the D largest singular values. \tilde{B}^{σ_L} is still right-normalized. The next $A^{\sigma_{L-1}}$ to the left, \tilde{U} and \tilde{S} are multiplied together to form $M^{\sigma_{L-1}}$. By reshaping and singular value decomposing,

$$M_{i,j}^{\sigma} = M_{i,(\sigma,j)} = \sum_k U_{i,k} S_k B_{k,(\sigma,j)} = \sum_k U_{i,k} S_k B_{k,j}^{\sigma}, \tag{2.21}$$

we obtain right-normalized $B^{\sigma_{L-1}}$, truncate U , S and $B^{\sigma_{L-1}}$ to \tilde{U} , \tilde{S} and $\tilde{B}^{\sigma_{L-1}}$, and the procedure continues. At the end, the compressed MPS $|\tilde{\psi}\rangle$ is right-normalized and given by \tilde{B} -matrices.

The disadvantage of the procedure is that there is a one-sided dependence of truncations on each other: truncations further to the left depend on those to the right, but not vice versa. If the truncation is small, the introduced additional inaccuracy is minor; however, for cases where large truncations may occur, the dependence might become too strong and the truncation far from optimal.

The optimal approach is to start from an ansatz MPS of the desired reduced dimension, and to minimize its distance to the MPS to be approximated iteratively, i.e. by changing its M^{σ} matrices iteratively. The matrices play the role of variational parameters.

Minimizing $\| |\psi\rangle - |\tilde{\psi}\rangle \|_2^2$ with respect to $|\tilde{\psi}\rangle$ is a highly nonlinear optimization problem in the \tilde{M}^{σ_i} . But this can be done iteratively and linearly as follows. Starting with an initial guess for $|\tilde{\psi}\rangle$, we sweep through the set of \tilde{M}^{σ_i} site by site, keeping all other matrices fixed and choosing the new \tilde{M}^{σ_i} , such that the distance is minimized. Repeating this sweep through the matrices several times will lead to a converged optimal approximation.

The new \tilde{M}^{σ_i} is found by extremizing with respect to $\tilde{M}_{a_{i-1},a_i}^{\sigma_i*}$, which only shows up in $-\langle \tilde{\psi} | \psi \rangle + \langle \tilde{\psi} | \tilde{\psi} \rangle$. We find

$$\begin{aligned} & \frac{\partial}{\partial \tilde{M}_{a_{i-1},a_i}^{\sigma_i*}} (\langle \tilde{\psi} | \tilde{\psi} \rangle - \langle \tilde{\psi} | \psi \rangle) \\ &= \sum_{\sigma^*} (\tilde{M}^{\sigma_1^*} \dots \tilde{M}^{\sigma_{i-1}^*})_{1,a_{i-1}} (\tilde{M}^{\sigma_{i+1}^*} \dots \tilde{M}^{\sigma_L^*})_{a_i,1} \tilde{M}^{\sigma_1} \dots \tilde{M}^{\sigma_i} \dots \tilde{M}^{\sigma_L} \\ & - \sum_{\sigma^*} (\tilde{M}^{\sigma_1^*} \dots \tilde{M}^{\sigma_{i-1}^*})_{1,a_{i-1}} (\tilde{M}^{\sigma_{i+1}^*} \dots \tilde{M}^{\sigma_L^*})_{a_i,1} M^{\sigma_1} \dots M^{\sigma_i} \dots M^{\sigma_L} = 0. \end{aligned}$$

The sum over σ^* runs over all physical sites except i . Assuming that we keep $|\tilde{\psi}\rangle$ at each iteration in mixed-canonical form, $\tilde{A}^{\sigma_1} \dots \tilde{A}^{\sigma_{i-1}} \tilde{M}^{\sigma_i} \tilde{B}^{\sigma_{i+1}} \dots \tilde{B}^{\sigma_L}$, the first term simplifies owing to the normalization conditions to \tilde{M}^{σ_i} . Then, we obtain the equation

$$\tilde{M}_{a_{i-1},a_i}^{\sigma_i} = O_{a_{i-1},a_i}^{\sigma_i}, \tag{2.22}$$




Figure 3. Equation for iterative compression of an MPS for a suitably normalized state. The fatter lines correspond to the state to be compressed, the thinner lines to the compressed state.

where $O_{a_{i-1}, a_i}^{\sigma_i}$ is given by

$$\sum_{a'_{i-1} a'_i} L_{a_{i-1}, a'_{i-1}} M_{a'_{i-1}, a'_i}^{\sigma_i} R_{a'_i, a_i}, \tag{2.23}$$

with L and R as indicated in figure 3. In fact, calculating L and R is nothing but carrying out the first steps of an overlap calculation, starting from the left or right. If one sweeps through the system from left to right and back, one can build L and R iteratively from previous steps, which is the most efficient way.

To make this work for an entire chain, we have to shift the boundary between the left- and right-normalized matrices as we move through the chain. But this can be done by applying one step of normalization (by SVD) to the matrices on a single site.

(b) Matrix product operators

In the previous section, we have considered a local operator $\hat{O}^{\sigma_i, \sigma'_i}$, which ties in nicely with the local MPS notation. For operators like the Hamiltonian \hat{H} , which consists of sums of (in practice usually) local terms of operator products, the natural question arises whether such operators can also be expressed in a form resembling an MPS. This is indeed the case and leads directly to MPOs. The most general operator acting on our L -site system is given as

$$\hat{O} = \sum_{\sigma, \sigma'} c_{\sigma, \sigma'} |\sigma\rangle \langle \sigma'|. \tag{2.24}$$

Interpreting the scalar coefficients as $c_{(\sigma_1, \sigma'_1), \dots, (\sigma_L, \sigma'_L)}$, we can, by analogy to an MPS, conclude that SVDs allow a decomposition as

$$\hat{O} = \sum_{\sigma, \sigma'} W^{\sigma_1 \sigma'_1} W^{\sigma_2 \sigma'_2} \dots W^{\sigma_L \sigma'_L} |\sigma\rangle \langle \sigma'|, \tag{2.25}$$

where at each lattice site, we introduce a set of d^2 matrices $W^{\sigma_i \sigma'_i}$. The previous statements about gauge degrees of freedom, orthonormalization types, etc. generalize directly to MPOs.

The graphical representation of an MPO is also an obvious extension from the MPS case, with an ingoing physical leg down and outgoing physical leg going up (figure 4).




Figure 4. Graphical representation of a matrix product operator (MPO), with in- and outgoing physical legs sticking out horizontally.




Figure 5. Graphical representation of the application of an MPO to an MPS. Matching legs are contracted over, the resulting MPS multiplies the dimensions of the underlying objects.

The application of an MPO to an MPS is very elegant (figure 5),

$$\begin{aligned}
 \hat{O}|\psi\rangle &= \sum_{\sigma, \sigma'} (W^{\sigma_1, \sigma'_1} W^{\sigma_2, \sigma'_2} \dots) (M^{\sigma'_1} M^{\sigma'_2} \dots) |\sigma\rangle \\
 &= \sum_{\sigma, \sigma'} \sum_{a, b} (W_{1, b_1}^{\sigma_1, \sigma'_1} W_{b_1, b_2}^{\sigma_2, \sigma'_2} \dots) (M_{1, a_1}^{\sigma'_1} M_{a_1, a_2}^{\sigma'_2} \dots) |\sigma\rangle \\
 &= \sum_{\sigma, \sigma'} \sum_{a, b} (W_{1, b_1}^{\sigma_1, \sigma'_1} M_{1, a_1}^{\sigma'_1}) (W_{b_1, b_2}^{\sigma_2, \sigma'_2} M_{a_1, a_2}^{\sigma'_2}) \dots |\sigma\rangle \\
 &= \sum_{\sigma} \sum_{a, b} N_{(1,1), (b_1, a_1)}^{\sigma_1} N_{(b_1, a_1), (b_2, a_2)}^{\sigma_2} \dots |\sigma\rangle \\
 &= \sum_{\sigma} N^{\sigma_1} N^{\sigma_2} \dots |\sigma\rangle.
 \end{aligned}$$

The elegance stems from the observation that the form of the MPS remains invariant, with an increase of the matrix size: the new MPS matrices,

$$N_{(b_{i-1}, a_{i-1}), (b_i, a_i)}^{\sigma_i} = \sum_{\sigma'_i} W_{b_{i-1} b_i}^{\sigma_i \sigma'_i} M_{a_{i-1} a_i}^{\sigma'_i}, \tag{2.26}$$

have the multiplied dimension of the MPS and the MPO. At the same time, like in the case of an overlap, the operation is only of small polynomial, not exponential complexity. Of course, it remains to show that it is in fact possible to construct

the MPO of, e.g. an Hamiltonian ‘by hand’, because the conceptual way of using an SVD might be exponentially complicated. But as we will see, this is no big problem indeed [13,14].

Let us consider the following simple Hamiltonian

$$\hat{H} = \sum_{i=1}^{L-1} \left\{ \frac{J}{2} \hat{S}_i^+ \hat{S}_{i+1}^- + \frac{J}{2} \hat{S}_i^- \hat{S}_{i+1}^+ + J^z \hat{S}_i^z \hat{S}_{i+1}^z \right\}. \tag{2.27}$$

As this is a shorthand for sums of tensor products of operators while omitting all the identity operators on other sites, it is convenient to reconsider the building block $W_{bb'}^{\sigma\sigma'}$ combined with its associated projector $|\sigma\rangle\langle\sigma'|$ to become an operator-valued matrix $\hat{W}_{b,b'} = \sum_{\sigma\sigma'} W_{b,b'}^{\sigma\sigma'} |\sigma\rangle\langle\sigma'|$ such that the MPO takes the simple form

$$\hat{O} = \hat{W}^{[1]} \hat{W}^{[2]} \dots \hat{W}^{[L]}. \tag{2.28}$$

In order to construct the MPO, we move through an arbitrary operator ‘string’ appearing in \hat{H} : starting from the right end, the string contains unit operators, until at one point, we encounter one of (in our example) three non-trivial operators. For the field operator, the string part further to the left may only contain unit operators; for the interaction operators, the complementary operator must follow immediately to complete the interaction term, to be continued by unit operators further to the left. For book-keeping, we introduce five corresponding states of the string at some given bond: state 1, only units to the right; states 2,3,4, one \hat{S}^+ , \hat{S}^- , \hat{S}^z just to the right; state 5, completed interaction somewhere to the right. Comparing the state of a string left and right of one site, only a few transitions are allowed: $1 \rightarrow 1$ by the unit operator \hat{I} , $1 \rightarrow 2$ by \hat{S}^+ , $1 \rightarrow 3$ by \hat{S}^- , $1 \rightarrow 4$ by \hat{S}^z . Furthermore $2 \rightarrow 5$ by $(J/2)\hat{S}^-$, $3 \rightarrow 5$ by $(J/2)\hat{S}^+$ and $4 \rightarrow 5$ by $J^z\hat{S}^z$, to complete the interaction term, and $5 \rightarrow 5$ for a completed interaction by the unit operator \hat{I} . Furthermore, all string states must start at 1 to the right of the last site and end at 5 (i.e. the dimension D_W of the MPO to be) to the left of the first site. This can now be encoded by the following operator-valued matrices:

$$\hat{W}^{[i]} = \begin{bmatrix} \hat{I} & 0 & 0 & 0 & 0 \\ \hat{S}^+ & 0 & 0 & 0 & 0 \\ \hat{S}^- & 0 & 0 & 0 & 0 \\ \hat{S}^z & 0 & 0 & 0 & 0 \\ 0 & (J/2)\hat{S}^- & (J/2)\hat{S}^+ & J^z\hat{S}^z & \hat{I} \end{bmatrix}, \tag{2.29}$$

and on the first and last sites

$$\hat{W}^{[1]} = [0 \quad (J/2)\hat{S}^- \quad (J/2)\hat{S}^+ \quad J^z\hat{S}^z \quad \hat{I}] \quad \hat{W}^{[L]} = \begin{bmatrix} \hat{I} \\ \hat{S}^+ \\ \hat{S}^- \\ \hat{S}^z \\ 0 \end{bmatrix}. \tag{2.30}$$

Inserting the explicit operator representations gives the desired $W^{\sigma\sigma'}$ -matrices for the MPO. The MPO in this example has dimension $D_W = 5$. For longer-ranged Hamiltonians, further ‘intermediate states’ have to be introduced: for

$$\hat{H} = J_1 \sum_i \hat{S}_i^z \hat{S}_{i+1}^z + J_2 \sum_i \hat{S}_i^z \hat{S}_{i+2}^z, \quad (2.31)$$

the bulk operator would read

$$\hat{W}^{[i]} = \begin{bmatrix} \hat{I} & 0 & 0 & 0 \\ \hat{S}^z & 0 & 0 & 0 \\ 0 & \hat{I} & 0 & 0 \\ 0 & J_1 \hat{S}^z & J_2 \hat{S}^z & \hat{I} \end{bmatrix}, \quad (2.32)$$

as one can verify by working out the operator string paths allowed by this construction. Except for special cases like exponentially decaying interactions, which find a very compact representation [15,16], the dimension of the MPO will grow with interaction range; however, we may always compress an MPO like an MPS to a smaller dimensional object with a minimal loss of information.

(c) *The finite-system density-matrix renormalization group algorithm*

The so-called *finite-system DMRG algorithm* is nothing but a variational ground-state search in the ansatz space of D -dimensional MPSs; hence, it is also called vMPS. Let me mention that historically, a finite-system DMRG was not framed in the language of an MPS, but rather in terms of the analysis of reduced density operators of subsystems (blocks) of the chain. If, on the other hand, one assumes that MPSs are very efficient at encoding one-dimensional states (as they are, see §2e), then the following algorithm for ground-state searches in that state class follows naturally—and it then turns out that White [1] had had the perfect intuition, as the algorithms are identical.

As increasing D leads to ansatz spaces that are supersets of the previous ones, the quality of approximation improves monotonically; in view of the exact MPS representation of quantum states, the limit $D \rightarrow \infty$ is exact. These two features allow for reliable extrapolation of results obtained for sequences of finite D .

Finding the ground state is equivalent to minimizing

$$E[|\psi\rangle] = \frac{\langle \psi | \hat{H} | \psi \rangle}{\langle \psi | \psi \rangle} \quad (2.33)$$

or, using a Lagrange multiplier λ , minimizing

$$\langle \psi | \hat{H} | \psi \rangle - \lambda \langle \psi | \psi \rangle. \quad (2.34)$$

Inserting the MPS encoding of $|\psi\rangle$, this turns into a highly nonlinear problem in the variables $M_{a_{i-1}, a_i}^{\sigma_i}$ which is unsolvable in numerical practice. But in fact, this problem can be turned into a sequence of linear problems, whose solutions will lead to an iterative improvement of the solution (in the sense that energy is monotonically lowered), where, with suitable numerical tricks, it can be ensured (in the sense of actual numerical practice) that the (energetically) optimal state in this ansatz class can be reached. One starts from a guess state, which can be




Figure 6. Standard eigenvalue problem for the optimization of $M_{a_{\ell-1}, a_{\ell}}^{\sigma_{\ell}}$. The unknown matrix is circled on the left network. The structure to the left of it is referred to as L in the text, the right structure as R .

picked either randomly (i.e. A^{σ} matrices with random entries) or be obtained from some warm-up algorithm such as the so-called *infinite-system* DMRG algorithm [1,2]. We consider the MPS state provided by the initial guess and pick one site ℓ , while all other sites remain inert. We then optimize energy with respect to the $A^{\sigma_{\ell}}$ matrices belonging to this site. Then, the variables appear in equation (2.34) only in a quadratic form, for which the determination of the extremum is a benign linear algebra problem. This will lower the energy, and find a variationally better state, but of course not the optimal one. Now one continues to vary the matrix elements on another site for finding a state again lower in energy, moving through all sites multiple times, until the energy does not improve anymore.

Let us analyse the parts of equation (2.34). If we keep the state suitably mixed normalized, the normalization conditions imply that

$$\langle \psi | \psi \rangle = \sum_{\sigma_{\ell}} \sum_{a_{\ell-1} a_{\ell}} M_{a_{\ell-1}, a_{\ell}}^{\sigma_{\ell}*} M_{a_{\ell-1}, a_{\ell}}^{\sigma_{\ell}}. \tag{2.35}$$

Considering $\langle \psi | \hat{H} | \psi \rangle$, with \hat{H} in MPO language and taking the extremum of equation (2.34) with respect to $M_{a_{\ell-1}, a_{\ell}}^{\sigma_{\ell}*}$, we find

$$\sum_{\sigma'_{\ell}} \sum_{a'_{\ell-1} a'_{\ell}} \sum_{b_{\ell-1}, b_{\ell}} L_{b_{\ell-1}}^{a_{\ell-1}, a'_{\ell-1}} W_{b_{\ell-1}, b_{\ell}}^{\sigma_{\ell}, \sigma'_{\ell}} R_{b_{\ell}}^{a_{\ell}, a'_{\ell}} M_{a'_{\ell-1}, a'_{\ell}}^{\sigma'_{\ell}} - \lambda M_{a_{\ell-1}, a_{\ell}}^{\sigma_{\ell}} = 0, \tag{2.36}$$

with the mathematical objects as defined graphically in figure 6. This is in fact a very simple eigenvalue equation; if we introduce a matrix H by reshaping $H_{(\sigma_{\ell} a_{\ell-1} a_{\ell}), (\sigma'_{\ell} a'_{\ell-1} a'_{\ell})} = \sum_{b_{\ell-1}, b_{\ell}} L_{b_{\ell-1}}^{a_{\ell-1}, a'_{\ell-1}} W_{b_{\ell-1}, b_{\ell}}^{\sigma_{\ell}, \sigma'_{\ell}} R_{b_{\ell}}^{a_{\ell}, a'_{\ell}}$ and a vector v with $v_{\sigma_{\ell} a_{\ell-1} a_{\ell}} = M_{a_{\ell-1}, a_{\ell}}^{\sigma_{\ell}}$, we arrive at an eigenvalue problem of matrix dimension $(dD^2 \times dD^2)$,

$$Hv - \lambda v = 0. \tag{2.37}$$

Solving for the lowest eigenvalue λ_0 gives us a $v_{\sigma_{\ell} a_{\ell-1} a_{\ell}}^0$, which is reshaped back to $M_{a_{\ell-1}, a_{\ell}}^{\sigma_{\ell}}$, λ_0 being the current ground-state energy estimate. In general, dD^2 is too large for an exact diagonalization, but as we are only interested in the lowest eigenvalue and eigenstate, an iterative eigensolver that aims for the ends of the spectrum will do. Typical methods are the Lanczos or Jacobi–Davidson large sparse matrix solvers.

The optimal algorithm then runs as follows.

- Start from some initial guess for $|\psi\rangle$, which is right-normalized, i.e. consists of B -matrices only.
- Calculate the R -expressions iteratively for all site positions $L - 1$ through to 1 iteratively.
- *Right sweep*: starting from site $\ell = 1$ through to site $L - 1$, sweep through the lattice to the right as follows: solve the standard eigenproblem by an iterative eigensolver for M^{σ_ℓ} , taking its current value as the starting point. Once the solution is obtained, left-normalize M^{σ_ℓ} into A^{σ_ℓ} by an SVD to maintain the desired normalization structure. The remaining matrices of the SVD are multiplied to the $M^{\sigma_{\ell+1}}$ to the right, which will be the starting guess for the eigensolver for the next site. Build iteratively the L expression by adding one more site. Move on by one site, $\ell \rightarrow \ell + 1$, and repeat.
- *Left sweep*: starting from site $\ell = L$ through to site 2, sweep through the lattice to the left as follows: solve the standard eigenproblem by an iterative eigensolver for M^{σ_ℓ} , taking its current value as the starting point. Once the solution is obtained, right-normalize M^{σ_ℓ} into B^{σ_ℓ} by an SVD to maintain the desired normalization structure. The remaining matrices of the SVD are multiplied to the $M^{\sigma_{\ell-1}}$ to the left, which will be the starting guess for the eigensolver for the next site. Build iteratively the R expression by adding one more site. Move on by one site, $\ell \rightarrow \ell - 1$, and repeat.
- Repeat right and left sweeps, until convergence is achieved. Convergence is achieved if energy converges, but the best test is (using MPOs) to consider $\langle \psi | \hat{H}^2 | \psi \rangle - (\langle \psi | \hat{H} | \psi \rangle)^2$ to see whether an eigenstate has been reached; this expression should approach 0 as closely as possible.

In this iterative process, the energy can only go down, as we continuously improve by varying the parameters. In fact, this most simple version of the DMRG can be improved in terms of stability and performance but the core algorithm remains unchanged; for details, see e.g. Schollwöck [5].

(d) *Variational or renormalization group method?*

In the exposition just given, the DMRG clearly emerges as a variational method, and the link to true RG methods (with scaling and truncating) is quite tenuous. It can, however, be put on a more solid footing if one considers the reduced density operators of subsystems and takes those to the thermodynamic limit [6]. For these, a fixed point relationship holds. The DMRG, therefore, sets up an RG flow in the space of reduced density operators, not in the space of Hamiltonians, as other RG methods would do.

(e) *Why does it work and why does it fail?*

The performance of all methods presented here rests on whether a quantum state can be approximated well by an MPS with bonds of a manageable dimension D . As we can cut a state into two parts across such a bond, we have to

consider bipartite quantum systems AB . Any state $|\psi\rangle = \sum_{ij} \psi_{ij} |i\rangle_A |j\rangle_B$, where the states $|i\rangle_A$ and $|j\rangle_B$ form orthonormal bases of dimensions N_A and N_B , respectively. Thinking of the ψ_{ij} as entries of a rectangular matrix Ψ (dimension $N_A \times N_B$), the reduced density matrices ρ_A and ρ_B take the forms

$$\rho_A = \Psi \Psi^\dagger \quad \text{and} \quad \rho_B = \Psi^\dagger \Psi. \quad (2.38)$$

If we assume $|\psi\rangle$ to be normalized, the eigenvalues w_a of ρ_A are non-negative, and sum to 1. This allows us to interpret them directly as statistical weights. If we assume that we know $|\psi\rangle$ exactly, the question of approximability reduces to how much statistical weight rests in the D eigenstates with the largest eigenvalues. In the few cases where this is the case, such analyses have been carried out in one and two dimensions starting with Peschel *et al.* [17] and Okunishi *et al.* [18]. They reveal that in one dimension for gapped systems, eigenvalues w_a generically decay exponentially fast (roughly as $e^{-c \ln^2 a}$), which explains the success of the DMRG, but in two-dimensional stripe geometries of size $L \times W$, $L \gg W$, $c \propto 1/W$, such that with increasing width W (increasing two dimensionality) the eigenspectrum decay is so slow as to make DMRG inefficient.

Usually, we have no clear idea about the eigenvalue spectrum; but it turns out that in such cases, entanglement entropies can serve as ‘proxy’ quantities, namely the von Neumann entanglement or entanglement entropy. It is given by the non-vanishing part of the eigenvalue spectrum of ρ_A as

$$S_{A|B} = -\text{Tr} \rho_A \log_2 \rho_A = - \sum_{\alpha} w_{\alpha} \log_2 w_{\alpha}. \quad (2.39)$$

If we now consider a bipartitioning $A|B$ where AB is in the thermodynamic limit and A of size L^D , with D the spatial dimension, the so-called *area laws* [19–23] predict that for ground states of short-ranged Hamiltonians with a gap to excitations entanglement entropy are not extensive, but proportional to the surface, i.e. $S(A|B) \sim L^{D-1}$, as opposed to thermal entropy. This implies $S \sim \text{cst.}$ in one dimension and $S \sim L$ in two dimensions. At criticality, a much richer structure emerges, which usually involves the presence or absence of logarithmic corrections ($S(L) \propto \log L$, $S(L) \propto L \log L$ in one dimension and two dimensions), see Srednicki [20], Vidal *et al.* [24], Latorre *et al.* [25], Goev & Klich [26] and Barthel *et al.* [27]. It should be emphasized that these properties of ground states are highly unusual: in the thermodynamic limit, a random state out of Hilbert space will indeed show extensive entanglement entropy with probability 1.

In a mathematically non-rigorous way, one can now make contact between the DMRG and the area laws of quantum entanglement: between two D -dimensional state spaces for A and B , as provided by a D -dimensional MPS, the maximal entanglement is $\log_2 D$ in the case where all eigenvalues of ρ_A are identical and D^{-1} (such that ρ_A is maximally mixed); meaning that one needs a state of dimension 2^S and more to encode entanglement S properly. This implies that for gapped systems in one dimension, an increase in system size will not lead to a strong increase in D ; in two dimensions, $D \sim 2^L$, such that the DMRG will fail, even for relatively small system sizes, as resources, i.e. MPS dimensions, have to grow exponentially (this however does not exclude very precise results for small

two-dimensional clusters or quite large stripes). Critical systems in one dimension are borderline cases: the logarithmic correction makes them untreatable in the thermodynamic limit, but enough data for excellent finite-size extrapolations can be obtained.

3. Generalizations of the density-matrix renormalization group ansatz to higher dimensions

(a) Tensor network states

Extensive attempts to simulate two-dimensional quantum systems within the conventional paradigm of DMRG as a one-dimensional method have been made over the last 15 years; generically, the two-dimensional system is seen as a one-dimensional snake where interactions that on the true lattice are short-ranged are now fictitiously long-ranged (with a range corresponding to the smaller of the two linear sizes). As discussed above, the entanglement scaling implies a matrix dimension that grows exponentially with system size; nevertheless, impressive results on wide strips have been obtained ([28,29] among the first and most recent ones), but it remains clear that except for astute finite-size extrapolations, there is no way towards the thermodynamic limit.

On the other hand, we may interpret the two matrix indices of M^{σ_i} as bond indices, with the row and column corresponding to bonds going to the left and right neighbour on a one-dimensional chain, as already implied by the graphical representation. It is, therefore, a natural extension to consider *tensors* $T_{\text{udlr}}^{\sigma_i}$ with four indices that correspond to the four bonds up, down, left and right on a two-dimensional square lattice and similar generalizations on other lattices (triangular, hexagonal and so on). Instead of an MPS, we would consider [10] a tensor product state (TPS) or tensor network states (TNSs)—a popular other name being projected entangled pair states (PEPS; [11])—reading

$$|\psi\rangle = \sum_{\sigma} \mathcal{C}[T^{\sigma_1}, T^{\sigma_2}, \dots, T^{\sigma_L}]|\sigma\rangle, \quad (3.1)$$

where the \mathcal{C} stands for a contraction over all connected legs in a tensor network; in the one-dimensional case this reduces immediately to the matrix multiplication of an MPS. For a graphical representation, see figure 7.

The attractive feature of this ansatz is the following: imagine a vertical cut through a system of size $L \times L$. In a PEPS, this will cut L legs of dimension D each. This allows for D^L configurations. The maximum amount of entanglement across the cut that can be encoded is given by the logarithm as $L \log D$; this means that it is proportional to the surface and hence the area law for gapped two-dimensional systems is immediately obeyed, indicating that this is a valid ansatz class in two dimensions.

It should be remarked that the concept of a TNS is much more general than just a set of tensors reflecting the underlying lattice structure. The TNS can also encode hierarchy of scales and make use of disentanglers to keep descriptions as efficient as possible, as is done by the multi-scale entanglement renormalization ansatz (MERA) scheme [30].




Figure 7. Two-dimensional PEPS tensor network state on a square lattice; each blob corresponds to a physical site. A bipartition along the dashed line cuts L legs of dimension D each, leading to the observance of an area law. The network can also be interpreted as the successive application of an MPO (in dotted box) to an MPS.

(b) *Tensor network algorithms: finding a projected entangled pair states ground state*

At first sight, the solution of two-dimensional quantum ground-state problems seems immediately at hand. We may also generalize from MPOs to tensor product operators (TPOs), express a Hamiltonian as a TPO efficiently, and solve the variational ground-state equation iteratively as in one dimension, varying one tensor after another while keeping all others fixed.

The problems that occur are twofold: (i) the contraction of the tensor network can only be approximate, never exact as in the one-dimensional case and (ii) numerical stability is an issue as orthonormalization properties are lost. The first is immediately obvious if one tries to contract two tensors versus two matrices: the product of two matrices is again a matrix, $\sum_j M_{ij} M_{jk} = M_{ik}$; however, the contraction of two 4-leg tensors over one leg produces one 6-leg tensor. This can be turned into an exact mathematical statement that the contraction of a two-dimensional tensor network is indeed a non-deterministic polynomial time hard problem, i.e. impossible on a classical computer. Approximation schemes, therefore, have to be involved. One example would be to read a PEPS as a succession of rows of tensors applied to an MPS at the bottom end, to be completed by an MPS at the top end. The tensor rows can be interpreted as MPOs and we would then have a succession of MPO–MPS contractions; these lead to a multiplicative growth of bond dimensions, which is checked by an MPS compression algorithm. But this is obviously only approximate. The second problem is less fundamental, but as annoying. In our exposition

of one-dimensional DMRG, we have made extensive use of the left- and right-normalizability of MPSs. But this was only possible because we considered open boundary conditions; for a ring (periodic boundary conditions), there is no separate left- and right-normalization. One immediate consequence is that the eigenproblem of the DMRG turns into a generalized eigenproblem, $Hv = \lambda Nv$, whose numerical stability depends on the condition number of N , which can turn out to be large. In a two-dimensional system, even with open boundary conditions, any tensor singled out for variational optimization generates a ring-like structure of tensors, implying potential numerical instability.

Let me conclude this outlook on two dimensions by a general statement on TNSs. As all approximative methods, they have some built-in bias. In the case of MPSs and TNSs, it is their preference for *low-entanglement* states. If the accuracy is not very high, and the energy resolution therefore not very good, among very close-lying low-energy states those with less entanglement will be preferred, which might bias to the wrong kind of ground state, e.g. a fixed singlet pattern over spin-liquid-like fluctuating ground states. In one-dimensional MPS (DMRG) applications, D can nowadays be pushed to sizes of 10 000 or so, resolving energies down to machine precision, and this bias does not seem to create serious problems. In two-dimensional applications, where current dimensions are very small, this issue is potentially much more severe!

4. Outlook

Another important extension of the DMRG I have completely omitted is the extension to the time domain, i.e. the simulation of quantum dynamics close and far from equilibrium. There are a variety of methods, all based on Vidal's insight that time evolutions of MPSs can be done very efficiently [31–35], which have given rise to numerous successful applications. Here, the current effort is mainly to try to extend the range of such simulations. In ground-state calculations, in one dimension, essentially all algorithmic questions of interest have been settled satisfactorily; in two dimensions, a sober assessment would indicate that at the moment, the best PEPS results do not yet significantly exceed the results of other methods, and the road to higher accuracies and higher stability seems very steep, but probably not hopeless.

References

- 1 White, S. R. 1992 Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **69**, 2863–2866. (doi:10.1103/PhysRevLett.69.2863)
- 2 White, S. R. 1993 Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B* **48**, 10 345–10 356. (doi:10.1103/PhysRevB.48.10345)
- 3 Schollwöck, U. 2005 The density-matrix renormalization group. *Rev. Mod. Phys.* **77**, 259–315. (doi:10.1103/RevModPhys.77.259)
- 4 Hallberg, K. 2006 New trends in density matrix renormalization. *Adv. Phys.* **55**, 477–526. (doi:10.1080/00018730600766432)
- 5 Schollwöck, U. 2011 The density-matrix renormalization group method in the age of matrix product states. *Ann. Phys.* **326**, 96–192. (doi:10.1016/j.aop.2010.09.012)
- 6 Ostlund, S. & Rommer, S. 1995 Thermodynamic limit of density matrix renormalization. *Phys. Rev. Lett.* **75**, 3537–3540. (doi:10.1103/PhysRevLett.75.3537)

- 7 Dukelsky, J., Martín-Delgado, M. A., Nishino, T. & Sierra, G. 1998 Equivalence of the variational matrix product method and the density matrix renormalization group applied to spin chains. *Europhys. Lett.* **43**, 457. (doi:10.1209/epl/i1998-00381-x)
- 8 Takasaki, H., Hikihara, T. & Nishino, T. 1999 Fixed point of the finite system DMRG. *J. Phys. Soc. Jpn* **68**, 1537–1540. (doi:10.1143/JPSJ.68.1537)
- 9 Verstraete, F., Porras, D. & Cirac, J. I. 2004 DMRG and periodic boundary conditions: a quantum information perspective. *Phys. Rev. Lett.* **93**, 227205. (doi:10.1103/PhysRevLett.93.227205)
- 10 Nishino, T., Hieida, Y., Okunishi, K., Maeshima, N., Akutsu, Y. & Gendiar, A. 2001 Two-dimensional tensor product variational formulation. *Prog. Theor. Phys.* **105**, 409–417. (doi:10.1143/PTP.105.409)
- 11 Verstraete, F. & Cirac, J. I. 2004 Renormalization algorithms for quantum-many body systems in two and higher dimensions. (<http://arxiv.org/abs/cond-mat/0407066>).
- 12 Golub, G. H. & van Loan, C. F. 1996 *Matrix computations*. Baltimore, MD: Johns Hopkins University Press.
- 13 McCulloch, I. 2007 From density-matrix renormalization group to matrix product states. *J. Stat. Mech. Theor. Exp.* **2007**, P10014. (doi:10.1088/1742-5468/2007/10/P10014)
- 14 Crosswhite, G. M. & Bacon, D. 2008 Finite automata for caching in matrix product algorithms. *Phys. Rev. A* **78**, 012356. (doi:10.1103/PhysRevA.78.012356)
- 15 Crosswhite, G. M., Doherty, A. C. & Vidal, G. 2008 Applying matrix product operators to model systems with long-range interactions. *Phys. Rev. B* **78**, 035116. (doi:10.1103/PhysRevB.78.035116)
- 16 Fröwis, F., Nebendahl, V. & Dür, W. 2010 Tensor operators: constructions and applications for long-range interaction systems. *Phys. Rev. A* **81**, 062337. (doi:10.1103/PhysRevA.81.062337)
- 17 Peschel, I., Kaulke, M. & Legeza, Ö. 1999 Density-matrix spectra for integrable models. *Ann. Phys.* **8**, 153–164. (doi:10.1002/(SICI)1521-3889(199902)8:2<153::AID-ANDP153>3.0.CO;2-N)
- 18 Okunishi, K., Hieida, Y. & Akutsu, Y. 1999 Universal asymptotic eigenvalue distribution of density matrices and corner transfer matrices in the thermodynamic limit. *Phys. Rev. E* **59**, R6227–R6230. (doi:10.1103/PhysRevE.59.R6227)
- 19 Bekenstein, J. D. 1973 Black holes and entropy. *Phys. Rev. D* **7**, 2333–2346. (doi:10.1103/PhysRevD.7.2333)
- 20 Srednicki, M. 1993 Entropy and area. *Phys. Rev. Lett.* **71**, 666–669. (doi:10.1103/PhysRevLett.71.666)
- 21 Callan, C. G. & Wilczek, F. 1994 On geometric entropy. *Phys. Lett. B* **333**, 55–61. (doi:10.1016/0370-2693(94)91007-3)
- 22 Plenio, M. B., Eisert, J., Dreissig, J. & Cramer, M. 2005 Entropy, entanglement, and area: analytical results for harmonic lattice systems. *Phys. Rev. Lett.* **94**, 060503. (doi:10.1103/PhysRevLett.94.060503)
- 23 Eisert, J., Cramer, M. & Plenio, M. B. 2010 Area laws for the entanglement entropy. *Rev. Mod. Phys.* **82**, 277–306. (doi:10.1103/RevModPhys.82.277)
- 24 Vidal, G., Latorre, J. I., Rico, E. & Kitaev, A. 2003 Entanglement in quantum critical phenomena. *Phys. Rev. Lett.* **90**, 227902. (doi:10.1103/PhysRevLett.90.227902)
- 25 Latorre, J. I., Rico, E. & Vidal, G. 2004 Ground state entanglement in quantum spin chains. *Quant. Inf. Comp.* **4**, 48–92. (<http://arxiv.org/abs/quant-ph/0304098>)
- 26 Goev, D. & Klich, I. 2006 Entanglement entropy of fermions in any dimension and the Widom conjecture. *Phys. Rev. Lett.* **96**, 100503. (doi:10.1103/PhysRevLett.96.100503)
- 27 Barthel, T., Chung, M.-C. & Schollwöck, U. 2006 Entanglement scaling in critical two-dimensional fermionic and bosonic systems. *Phys. Rev. A* **74**, 022329. (doi:10.1103/PhysRevA.74.022329)
- 28 White, S. R. 1996 Spin gaps in a frustrated Heisenberg model for CaV_4O_9 . *Phys. Rev. Lett.* **77**, 3633–3636. (doi:10.1103/PhysRevLett.77.3633)
- 29 White, S. R. & Chernychev, A. L. 2007 Neel order in square and triangular lattice Heisenberg models. *Phys. Rev. Lett.* **99**, 127004. (doi:10.1103/PhysRevLett.99.127004)

- 30 Vidal, G. 2007 Entanglement renormalization. *Phys. Rev. Lett.* **99**, 220405. (doi:10.1103/PhysRevLett.99.220405)
- 31 Vidal, G. 2003 Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.* **91**, 147902. (doi:10.1103/PhysRevLett.91.147902)
- 32 Vidal, G. 2004 Efficient simulation of one-dimensional quantum many-body systems. *Phys. Rev. Lett.* **93**, 040502. (doi:10.1103/PhysRevLett.93.040502)
- 33 Daley, A. J., Kollath, C., Schollwöck, U. & Vidal, G. 2004 Time-dependent density-matrix renormalization-group using adaptive effective Hilbert spaces. *J. Stat. Mech. Theor. Exp.* P04005. (doi:10.1088/1742-5468/2004/04/P04005)
- 34 White, S. R. & Feiguin, A. E. 2004 Real time evolution using the density matrix renormalization group. *Phys. Rev. Lett.* **93**, 076401. (doi:10.1103/PhysRevLett.93.076401)
- 35 Verstraete, F., Garcia-Ripoll, J. J. & Cirac, J. I. 2004 Matrix product density operators: simulation of finite-temperature and dissipative systems. *Phys. Rev. Lett.* **93**, 207204. (doi:10.1103/PhysRevLett.93.207204)