

Tensor Network Algorithms: a Route Map

Mari Carmen Bañuls^{1,2}

¹*Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Str.1, Garching, Germany, D-85748; email: mari.banuls@mpq.mpg.de*

²*Munich Center for Quantum Science and Technology (MCQST), Schellingstr. 4, Munich, Germany, D-80799*

Tensor networks provide extremely powerful tools for the study of complex classical and quantum many-body problems. Over the last two decades, the increment in the number of techniques and applications has been relentless, and especially the last ten years have seen an explosion of new ideas and results that may be overwhelming for the newcomer. This short review introduces the basic ideas, the best established methods and some of the most significant algorithmic developments that are expanding the boundaries of the tensor network potential. The goal is to help the reader not only appreciate the many possibilities offered by tensor networks, but also find their way through state-of-the-art codes, their applicability and some avenues of ongoing progress.

CONTENTS

I. INTRODUCTION	1
II. BASIC CONCEPTS	2
A. Tensor network states	2
B. Fundamental primitives	3
1. Contracting TN	4
2. Tensor update	4
3. Truncation	4
C. Classic algorithms	4
1. Variational optimization of MPS	4
2. Evolving MPS: TEBD, tMPS	5
III. ADVANCED TNS METHODS	6
A. Higher dimensions	6
B. Symmetries	7
C. Fermions	8
D. Dynamics	9
E. Excitations	9
IV. FURTHER TN APPROACHES AND PERSPECTIVES	10
A. Network renormalization approaches	10
B. Connections to other techniques	10
V. OUTLOOK	11
ACKNOWLEDGMENTS	11
References	11
RELATED RESOURCES	13

I. INTRODUCTION

Numerical methods are an essential tool to tackle quantum many-body systems, most of which lack analytical solutions. For these problems, though, the dimensionality—and with it the computational complexity—grows exponentially with the system size.

This limits the applicability of exact numerical calculations and calls for the development of numerical methods that can efficiently deal with, at least, the most relevant physical questions. Introduced to the field in the nineties, tensor network (TN) techniques aim to cover this need and have become by now, together with exact diagonalization and Quantum Monte Carlo methods, a key instrument in the numerical study of quantum many-body problems.

TN have been discovered independently in different disciplines. First uncovered in statistical physics by Baxter [1], in the field of quantum many-body physics their ancestry can be traced back to the first valence bond solid (VBS) proposed by Affleck, Kennedy, Lieb and Tasaki [2] as exact ground state of a short-range spin chain—see [3] for a review and historical perspective. Klümper et al. [4] later extended the AKLT proposal for a larger set of models, and also introduced the term *matrix product* to designate these states. The construction was generalized and formalized mathematically by Fannes, Nachtergaele and Werner [5] in the finitely correlated states for infinite spin chains.

Around the same time the density matrix renormalization group (DMRG), a new algorithm proposed by White [6], was revealing an amazing power to capture the ground state of large quantum spin chains with only modest numerical effort. Shortly afterwards, Östlund and Rommer [7] identified the fixed point of the infinite DMRG algorithm with precisely such matrix product states, and Dukelsky et al. [8] pointed out the connection between DMRG and a variational search over these states. Furthermore Nishino and Okunishi [9, 10] unified DMRG with Baxter’s corner transfer matrix approach for two dimensional classical models. And these insights inspired further generalizations of the original algorithm [11].

DMRG was applied to multiple scenarios and fast became a method of choice to study the static properties of quantum spin systems in low spatial dimension [12, 13]. Yet a whole new perspective was gained thanks to quantum information concepts. Understanding in terms of entanglement the matrix product ansatz [14] and the




FIG. 1. Graphical representation of tensors: (a) Example of a TN formed by four tensors; when contracted, a 4-rank tensor is obtained; (b) graphical representation of a TNS in each of the main families, for a system of 16 sites (note: triangles are commonly used to indicate isometries).

DMRG algorithm [15], and reformulating the latter fully in terms of matrix product states (MPS) [16] opened up the possibilities for improvements and jumpstarted the *tensor network* field. In particular, algorithms for real time evolution [17–19] and finite temperature [20–22] with matrix product states, as well as a generalization to higher dimensions [23] were proposed soon afterwards, revealing the potential of the tensor network picture.

Nowadays, TN algorithms are among the standard numerical methods for strongly correlated low-dimensional quantum systems. Most commonly used are the original methods from the early 2000s, which continuously find new applications. But the TN language continues to be exploited to provide, not only deeper mathematical understanding of the ansatz [24], but also new numerical techniques.

The variety of TN applications that have bloomed over the last decade and produced state-of-the-art results is too vast to do justice to it in these pages. Thus, the focus of this article is the general framework of TN algorithms, with a stress on a few selected advances in the field that

A. Tensor network states

In particular, a *tensor network state* (TNS) encodes all coefficients (in a given basis) of a quantum many-body state in such a diagram, with as many open legs as constituents in the system. Each dangling leg corresponds to the (finite) *physical* dimension of one site, while contracted legs correspond to *virtual* or *bond dimensions*.

TNS families are defined by graphs with different connectivities. For the families of interest, the number of parameters, proportional to the number of tensors, grows polynomially with the system size. This represents a drastic reduction with respect to the exponentially large dimension of the Hilbert space. But the aim of TNS is to capture physical states, which happen to explore only a small fraction of all possible quantum states, mainly characterized by their low entanglement. In particular, ground and thermal equilibrium states of local Hamiltonians fulfill an entanglement area law [31]: the entanglement between a certain subsystem and the rest scales

are important for cutting-edge applications. The details of the algorithms are not explicitly shown; the interested readers are encouraged to refer to the many excellent reviews in the literature, such as [25–30], to name only a few.

II. BASIC CONCEPTS

A tensor, the basic object, is simply a multidimensional array. The graphical representation of TN, illustrated in fig 1, provides a practical language to describe their algorithms and properties. For instance, a k -rank tensor, an object with k indices, is depicted as a geometrical shape with k legs (e.g. a matrix would have two legs). A contraction between two tensors—such as a matrix-vector product—is represented by joining the contracted indices. In general, a tensor network is a set of such interconnected tensors, resulting in a rank determined by the number of open legs (see fig 1a).

with the size of the boundary between both parts (or with small corrections thereof), instead of with the size of the bulk of the subsystems, as is the case for most states in the Hilbert space. A rigorous proof of the area law scaling exists for gapped one-dimensional local Hamiltonians [32] and for thermal equilibrium states in any dimension [33], whereas critical ground states can display small (logarithmic) corrections [34, 35].

The following are the most widely used TNS families (see their diagrams in fig. 1b).

1. MPS have a one-dimensional structure, with one tensor per lattice site [36]. Each tensor owns one open index corresponding to the physical dimension of the site, and two virtual legs connected to the neighboring sites (for open boundaries, the edge tensors only connect to one neighbor). The tensor for site k has components $A[k]_{\alpha\beta}^i$, where i takes values over the physical dimension (typically denoted d_k), and α and β respectively take values over

the left and right virtual dimensions of the tensor (denoted D_l and D_r)—equivalently, each $A[k]^i$ is a $D_l \times D_r$ matrix. More explicitly, for a system of N sites, all with physical dimension d , the state can be written

$$|\Psi\rangle = \sum_{i_1, i_2, \dots, i_N=1}^d \text{tr} (A[1]^{i_1} A[2]^{i_2} \dots A[N]^{i_N}) |i_1 i_2 \dots i_N\rangle. \quad (1)$$

MPS satisfy an entanglement area law: the half-chain entanglement of an MPS with maximal bond dimension D is upper-bounded by $S = 2 \log D$. Furthermore, they hold exponentially decaying correlations, can be prepared and contracted efficiently, and essentially correspond to ground states of local one-dimensional gapped Hamiltonians [37].

2. PEPS are the natural generalization of MPS to arbitrary graphs, where they can be defined with one tensor (with a physical leg) per vertex and connections according to the graph edges [23]. They can be expressed analogously to Eq. 1, replacing the trace by a contraction over all connections. PEPS fulfill the area law in higher dimensions, and are much more complex objects than MPS. For instance, they cannot—in the general case—be prepared or contracted efficiently [38] and, even with small bond dimension, they can support critical correlations [39].
3. TTN correspond to tree graphs. Usually—but not always—they have physical indices in the leave nodes [40], connected to tensors with only virtual indices at higher levels (see fig. 1b), which can correspond to a renormalization direction [41]. As MPS, TTN are loop-free and can be contracted efficiently, but they violate the one-dimensional area law, and can hold power-law decaying correlations when averaging over spatial positions [42]. TTN can be used also for higher dimensional systems [43].
4. MERA implement a more complex renormalization of the physical degrees of freedom [44–46], in which layers of unitary transformations (called disentanglers) that remove short range correlations are alternated with layers of isometries that perform the renormalization step. This results in a TN with cycles in which, thanks to the unitarity properties of the tensors, local expectation values can be computed efficiently. Scale invariant MERA can describe quantum critical ground states in one dimension [47, 48], where they support logarithmic corrections to the area law. However, in two dimensions they are proven to be a subset of PEPS [49], and thus satisfy the area law. However, a generalization called branching MERA [50] exists that can support up to volume-law entanglement in more than one spatial dimension.

Any tensor network has a so-called gauge freedom, since inserting the product of a matrix and its inverse XX^{-1} in between any contracted pair of indices (i.e. in a connected leg) leaves the whole TN invariant, but allows redefining pairs of neighboring tensors. For loop-free TNS, in which cutting a bond splits the network in two, it is possible to define a canonical form, in which the basis for the virtual index is chosen to be the Schmidt basis for the bipartition corresponding to the bond, explicitly encoding the corresponding entanglement [36]. Besides being fundamental to characterize the properties of a TNS family, this canonical form gives rise to more stable and efficient numerical algorithms.

The families above can be defined for finite-size systems with site-dependent tensors, but it is also possible to consider directly the thermodynamic limit, in which one (or a few) tensors are repeated infinitely many times, to produce a translationally invariant (or periodic in space) structure. In the case of MPS, the translationally invariant ansatz is called uniform MPS (uMPS). In infinite PEPS (iPEPS), a periodic iteration of a finite unit cell is most commonly used in practice, while the translationally invariant version is fundamental for the formal results [24]. This allows targeting bulk properties directly, without finite-size extrapolations, or, in the case of MERA, capturing the scale invariance of critical systems [47, 48].

These families can also describe mixed states. The simplest approach is to postulate the TNS ansatz in a given tensor product basis of the vector space of operators, with simply doubled physical legs. In particular, in the MPS and PEPS cases the resulting structures are called MPO [20, 21, 51] and PEPO. But if the ansatz is to describe a physical state, it must be positive semidefinite, a global property that cannot be assessed at the level of the local tensors. An alternative is to consider the TNS describing a purification, i.e. a pure state of the system plus an ancilla, such that tracing out the latter results in the desired mixed state. In the MPS and PEPS case, this yields a locally purified form, a TN with the same structure, where local tensors have double physical indices, and internal structure granting positivity [20]. This is more restrictive and potentially less efficient than the generic ansatz [52], but can be used in practice in numerical algorithms.

B. Fundamental primitives

Virtually all TN algorithms rest on two basic blocks: contracting (part of) the tensor network, and locally updating the tensors. Together with the approximation of (parts of a) TN by tensors with truncated dimensions, they can be considered the fundamental primitives on which more or less sophisticated higher-level algorithms are built.

1. Contracting TN

A ubiquitous problem in TN algorithms is contracting a tensor network. This means explicitly evaluating the products and sums of tensor components indicated by the connections, to result in a tensor with dimensionality corresponding to the indices that remain open (see fig.1a). For instance, for classical statistical models, partition functions and expectation values of local observables can be written as *closed* TN (without open indices). For TNS representing quantum states, norms and local expectation values are also closed tensor networks, while reduced density matrices appear as smaller tensor networks with operator indices. Two aspects of the contraction affect the implementation and performance of the algorithm.

1. Contraction order. In general, the computational cost of contracting a series of tensors with each other depends on the order in which operations are applied. For the regular networks that appear in the most common TNS algorithms, the number of possibilities is small, and the optimal sequence (which minimizes the computational cost) is known. But in the general case, finding the optimal contraction order is a NP-complete problem, for which some heuristic algorithms exist [53, 54].
2. Computational cost. If a contraction order exists whose computational cost grows only polynomially with the size of the network, we say that the TN can be contracted exactly. Such is the case with TN that do not contain loops, for instance the networks corresponding to expectation values of multi-point correlators in MPS and TTN. Also for TNS with some unitary properties there are contractible quantities, for instance the norm or few-point correlators evaluated in MERA. The exact contraction of an arbitrary tensor network is however a $\#P$ -complete problem [38]. Thus, most algorithms involving TN in more than one dimension need to approximate the contractions, which is referred to as *truncation* (IIB3).

2. Tensor update

Many algorithms work by holding a TN description of the quantity of interest and iteratively improving it until some predefined level of convergence is attained. The improvement proceeds by local changes, or updates, in which one (or few) tensors are modified in order to optimize the relevant cost function. Typically, the latter depends on all tensors in the network but only one is allowed to vary in each update step, while keeping the others fixed, hence turning the problem into a local one. A related concept is thus the *environment of a tensor*, the part of a TN that is complementary to the tensor being

modified. This appears in the local cost function and needs to be evaluated by a (in many cases approximate) contraction, in order to determine the proper update for the local tensor.

3. Truncation

Truncating a TN means reducing (some of) the dimensions of its tensors, ideally in such way that the global result does not change. A truncation can be part of an approximation strategy (e.g. for a PEPS, see III A), where it is used to control the dimension of a partial TN contraction. In the context of quantum states represented as TNS, truncating typically means finding tensors to approximate the state within a given family. For instance after acting with an operation that, if applied exactly, would increase the tensor dimensions, as is often the case of time evolution or non-local operators (e.g. IIC 2). And, more generally, truncating may refer to approximating a certain state with a TNS of fixed bond dimension.

In any truncation step, a decision is made as to which degrees of freedom to keep and which to discard. In a TNS, the fixed bond dimension upper-bounds the amount of correlations the state can hold, and thus the truncation step in most algorithms can be related to entanglement properties.

C. Classic algorithms

Numerous TN algorithms have been introduced in the last years, yet there are a few well-established methods that are used to obtain state-of-the-art results in quantum many-body problems. Many of them (specially for one-dimensional problems) are available as open-source implementations (see Related Resources at the end), making it possible to benefit from the numerical power of TN methods without the need to dive into implementation details. Furthermore, they are not difficult to implement, and can be easily adapted to solve other problems, beyond the ones they were originally designed for. They constitute the true workhorses of TNS numerical results.

1. Variational optimization of MPS

One of the most powerful strategies in the TNS toolbox is the variational optimization of the ansatz with respect to a given cost function, the paradigmatic example being the DMRG algorithm [6]. This can be essentially understood as an application of the variational principle in which an ansatz for the ground state is obtained minimizing the energy for a quantum many-body Hamiltonian over the set of MPS with fixed bond dimension




FIG. 2. Graphical expression of the local problems solved by the classic algorithms: (a) variational optimization for a single tensor in DMRG; (b) update of the local pair of tensors in TEBD (diamond-shaped tensors represent the Schmidt values, explicit in the canonical form); (c) local optimization in tMPS.

D [15, 26],

$$|\Phi_{\text{GS}}^{(D)}\rangle = \underset{|\Psi_D\rangle}{\text{argmin}} \frac{\langle \Psi_D | H | \Psi_D \rangle}{\langle \Psi_D | \Psi_D \rangle}. \quad (2)$$

The problem is tackled in an iterative manner, a single tensor being minimized at each step while the rest are kept constant. [55] While not strictly necessary, the implementation of the original method is greatly simplified by writing the Hamiltonian as a MPO [16]. This can be done exactly for short-range one-dimensional Hamiltonians [51], and approximation schemes exist for long-range interactions [e.g. [56]]. In this form, the local cost function can be written as the ratio of two tensor networks (see fig. 2a) that can be contracted efficiently with a cost that, for N sites, only scales as $O(ND^3)$. The local problem has thus the form of a Rayleigh-Ritz quotient, and can be solved exactly using a standard eigensolver. The procedure is iterated, sequentially optimizing each tensor in the ansatz, and repeatedly sweeping back and forth over the whole chain until a predetermined convergence criterion (usually convergence of the energy value within certain precision) has been reached. Further gain in efficiency is possible if tensors are always kept in canonical form and intermediate calculations are stored in memory. Because the optimum of each local problem can be found exactly, the algorithm is guaranteed to lower the energy monotonically, and thus to converge (even though this might be to a local minimum).

The infinite DMRG (iDMRG) algorithm directly targets systems in the thermodynamic limit [26], and can also be expressed in similar terms. In that case, instead of sweeping back and forth, at each step a unit cell of tensors is inserted and optimized in the middle of the chain, and the procedure is iterated until a fixed point has been reached.

The most natural scenario for the algorithm is the search for the ground state of a local one-dimensional Hamiltonian. The method is extremely competitive even for critical systems (for which the MPS can only approximate the correlations), thanks to finite-size and finite-entanglement [57, 58] scaling, and has been successfully used for long-range interactions and problems in larger dimensions (see sec. III A). The efficiency and robustness of the method make it one of the most powerful numerical methods available to solve quantum many-body prob-

lems. Additionally, it can be applied to any variational optimization problem in which the cost function is expressed in terms of an effective Hamiltonian with MPO structure [e.g. [59]].

2. Evolving MPS: TEBD, tMPS

The Time Evolved Block Decimation (TEBD) algorithm [14, 17] is arguably the simplest to implement, yet one of the most versatile methods in the TN toolbox. The strategy was originally proposed for simulating the evolution of an MPS under a quantum circuit, which can be written as a sequence of two-body, nearest-neighbor unitary gates. Since each gate can increase the entanglement, its exact action on an MPS generally results in a larger bond dimension. Maintaining an efficient description of the state thus requires an approximation step that reduces (truncates) the bond dimension after the application of each gate. The TEBD strategy proceeds via a local update, involving only the directly affected tensors, and corresponds to minimizing the distance between the transformed and updated states under the condition that all the remaining tensors are kept invariant. Exploiting the canonical form of the MPS, this can be achieved by a singular value decomposition of a single tensor, obtained when contracting together the gate and the local MPS tensors, including their environment, which encodes the state of the rest of the system (see fig. 2b). In the TEBD truncation step, only singular values above a certain threshold are kept, and the discarded weight gives a measure of the error.

For a one-dimensional nearest-neighbour Hamiltonian, the time evolution operator can be approximated, using a Trotter-Suzuki expansion, as a sequence of such two-body gates of the form $\exp(-i\delta h_i)$, where h_i is a two-body term and δ a short time step. The method can thus be used to simulate the dynamics of an MPS with cost that scales as $O(D^3)$, for bond dimension D . The scheme can be adapted for other (finite-range) Hamiltonians, although the cost increases steeply with the interaction range.

As an alternative to the local truncation, it is also possible to vary all tensors in order to minimize the distance to the exact state after one or more gates [39]. In this strategy, called tMPS,[60] tensors are optimized sequen-

tially as in the variational method IIC1, by solving a local problem that, in this case, reduces to a system of linear equations, also with cost $O(D^3)$ (see fig. 2c). In this way once can apply onto the MPS vector any MPO operator, in particular, a step of the Trotterized time-evolution. The cost of such MPO representation also increases with the range of the Hamiltonian, but long range interactions can be treated with help of an approximation scheme [61].

These methods are very efficient and extraordinarily versatile. Starting from an arbitrary state, the ground state can be approached by imaginary (or Euclidean) time evolution, which effectively projects the state onto its lowest energy component, and can be applied with the same algorithm [17], only using non-unitary terms $\exp(\delta h_i)$. Also thermal equilibrium states can be approximated using this technique [20–22], by writing a purification of the Gibbs ensemble (namely the thermofield state) as the evolution of a maximally entangled initial state in imaginary time given by the inverse temperature, $|\Psi\rangle \propto e^{-\beta\hat{H}/2} \sum_n |n\rangle|n\rangle$ (where n label a basis of the system Hilbert space). And by treating the mixed state as a vector in operator space, the same basic method can be used to simulate real time evolution of open systems under master equations [20, 21]. Imaginary time evolution of pure states can also be used to produce a sample of minimally entangled typical thermal states (METTS) [62] that reproduce thermal properties. These are only a few examples: more generally, the tMPS strategy approximates the action of any linear operator written as an MPO onto an MPS. This allows reformulating most linear algebra algorithms as approximate versions in the framework of MPS [e.g. [63, 64]]. And TEBD and tMPS algorithms can also be applied to translationally invariant (or periodic) MPS, working directly in the thermodynamic limit [25, 65].

Even though the technique to treat all the scenarios named above is almost identical, the entanglement in each of them, and thus the performance of the method, widely differs. While thermal equilibrium states satisfy an area law [33, 66] and admit efficient TNS approximations for local Hamiltonians [67, 68], real time evolution of a far-from-equilibrium state can give rise to linear growth of entanglement, in which case approximating the resulting state with an MPS would require the bond dimension to grow exponentially with the total time [69–71]. For this reason, while MPS methods are extremely useful to study dynamics close to equilibrium, or for moderate times [72], they suffer a fundamental limitation for genuinely out-of-equilibrium scenarios.

III. ADVANCED TNS METHODS

Even though the algorithms described in the previous section can treat a large number of problems, some more advanced techniques, developed mostly in the last decade and not yet freely available, are necessary to fully exploit

the power of TNS.

A. Higher dimensions

Despite the resounding success of the one-dimensional applications of TNS, applications of higher-dimensional ansatzes remain much less common. However, in the last years, the situation has started to change, thanks to a number of algorithmic developments and an intense effort of the community.

Treating two- and higher-dimensional problems has always been a coveted target of these numerical methods and, since the early times of DMRG, the possibility was recognized of applying the technique to two-dimensional quantum [73] and three-dimensional classical problems [11]. MPS form a complete family, and can be used as an ansatz for any problem, in particular in larger dimensions. For two-dimensional quantum states, the MPS ansatz can be wrapped around the lattice. This is usually done in a zig-zag or snake form, but other choices are possible [74]. The resulting representation of the Hamiltonian as MPO is more expensive (since some short-range terms get mapped onto longer-range ones), and a larger bond dimension is required to reach the desired convergence: since cutting a single bond partitions the state in two, to accommodate the entanglement of a state that satisfies an area law the bond dimension needs to grow exponentially with one of the dimensions of the system. Highly accurate computations are still obtained from systems of limited size, often exploiting a long-cylinder geometry and careful finite size extrapolations [e.g. [75, 76]].

With built-in area law, PEPS are a more suitable TNS ansatz, which supports good approximations for equilibrium states of local Hamiltonians [32, 68], and allows variational and evolution strategies as described in section IIC [23, 77]. They can also be used directly in the thermodynamic limit, in which case they are called iPEPS, and are parametrized by a unit cell with a finite number (which could be as small as one) of tensors [78]. Nevertheless, numerical algorithms with PEPS are considerably more involved, and have higher computational cost in terms of the tensor dimensions.

For starters, contracting PEPS is, in contrast to the efficient contraction of MPS, a ($\#P$ -complete) hard computational problem [38]. Practical algorithms resort to approximate contractions, in which the two-dimensional network is approximated as a sequence of MPO-MPS contractions from the boundary [77–79], by a coarse-graining, or tensor renormalization [80, 81] (see sec. IV A), or, in the case of iPEPS, by a corner transfer matrix contraction [10, 82]. In these strategies there is a trade-off between the numerical cost and the accuracy of the contraction, important to determine the environment of a tensor or the expectation values of observables.




FIG. 3. Approximation of the environment tensor in PEPS: (a) environment (in solid colors) of a pair of tensors on which a nearest-neighbor gate is applied (group shown in lighter shade); (b) in the simple update, the environment tensor is approximated as a product (compare to the 1D case in fig. 2b); (c) a correlated approximation of the environment is required for the full update.

Because of this, methods have been developed that gain efficiency by allowing less precise environment estimations for tensor updates (fig. 3). The most efficient alternative uses a so-called *simple update* [80] where, in order to update a pair of tensors under the action of a two-site gate, the environment is approximated by a product of diagonal matrices acting on each link surrounding the pair, which play a role analogous to that of the Schmidt values in the TEBD procedure. Discarding the correlations in the environment can prevent the method from reaching the best PEPS with fixed bond dimension in the general case [79], but the algorithm is still popular, due to its efficiency and stability. A better update can be found with the more expensive *full update* [78, 83], which takes into account a more accurate correlated environment approximation.

These approaches may improve the efficiency of the updates, which can be particularly useful in the case of ground state search by imaginary time evolution, where the goal is a fixed point of the evolution. However, the evaluation of observables still needs to be as accurate as possible, in order to guarantee a variational result. This yields for most of PEPS algorithms a computational cost scaling as $O(D^{10})$.

Another difference between PEPS and MPS computations is the absence of a canonical form for the former. As a consequence, the effective norm term appearing, for instance, in the denominator of Eq. 2, cannot be reduced to the identity, and needs to be inverted to solve the local problems, which results in higher computational costs and loss of stability. The problem can be alleviated making use of the gauge freedom to optimize the condition of this effective matrix [84–86].

Despite the higher computational challenge, and the still ongoing development of more efficient strategies, PEPS already outperform MPS for two-dimensional problems of moderate size, as explicitly shown in [87] for Heisenberg and Hubbard models.

All in all, iPEPS has been the preferred ansatz to address ground states of two-dimensional quantum problems in this context, due to the possibility of directly addressing bulk properties. [88] Imaginary time evolution has been the method predominantly used until

very recently, due to the highly non-linear character of a variational approach in the line of IIC. However, in the last few years, new strategies have been introduced for a stable and efficient variational optimization of iPEPS [89, 90], which produces more accurate results. A further step has been the precise solution of critical systems, with the help of extrapolations in the correlation length of finite D states [91–93]. Plenty of impressive numerical results have been already obtained thanks to these advanced methods, among them the most accurate result for the Hubbard model [94], and the first studies of three-dimensional problems [95].

When the focus is not on ground states, and similar to the MPS case discussed in section IIC 2, the (real or imaginary) time evolution techniques allow addressing multiple problems, such as equilibrium states at finite temperature [96], steady states of open systems [97–99] and real time evolution [77, 96, 100].

An alternative direction has been the development and exploitation of restricted subsets of PEPS, with more favourable computational properties, that can be suitable ansatzes for particular problems. It is the case, for instance, of sequentially generated states [101], or the more general isometric PEPS [102], or of Gaussian fermionic PEPS [103].

And other TNS families that do not have an area law, or only a restricted one, can be used to study higher-dimensional systems of a certain size, as does two-dimensional DMRG. For instance TTN [e.g. [43, 104]], or the recently introduced augmented trees [105].

B. Symmetries

In case the problem under study exhibits some symmetry, taking advantage of it is not only of fundamental interest, but can also boost the performance of a numerical algorithm. For instance, if the Hamiltonian commutes with a certain operator $[H, O] = 0$, its eigenstates will have well-defined eigenvalues of O , and the search can be restricted to subspaces labelled by particular quantum numbers.

In the case of quantum many-body systems, one is of-

ten interested in problems with a global symmetry of the form $U^{\otimes N}|\Psi\rangle = |\Psi\rangle$, where U is a unitary transformation that acts on a single site. Particularly relevant is the case when the operation is a representation of a group G , namely $U = U_g$, for some $g \in G$. Such Abelian symmetries were soon incorporated to the DMRG method [16, 26], where they became of common use, typically implemented for the conservation of particle number or total magnetization. Also the formalism for non-Abelian symmetries was also developed [8, 106], albeit not so commonly used.

A general framework to handle global symmetries in higher dimensional TN was first introduced in [107], with explicit formulations for Abelian [108, 109] and non-Abelian [110, 111] cases following shortly. The basic idea of these and the original DMRG constructions is to define invariant tensors, which remain unchanged when the symmetry operation acts on all the indices. This requires well-defined transformation properties for each of the indices and, in particular, choosing bases for the virtual legs with well-defined quantum numbers, for instance $|q\alpha\rangle$, where q labels an irreducible representation of the group, [112] and α labels the states within the same irrep. The bond dimension of such leg will be the sum of dimensions for each q . Assigning a direction to each edge in the TN, outgoing and incoming indices transform respectively with the unitary representation of the group and its inverse and it follows that a TNS constructed out of such invariant tensors is globally invariant.

The invariance of a tensor implies some internal structure. In the case of Abelian symmetries, the tensor can be decomposed in a direct sum of blocks, with the only non-vanishing ones being those for which the sum of quantum numbers of incoming indices equals that of outgoing ones. In the non-Abelian case, blocks corresponding to a suitable combination of irreps have further structure, as they can be further decomposed as a tensor product of one part dictated merely by the symmetry, and another one containing the free parameters of the state. In particular, for three-legged tensors the first factor is a Clebsch-Gordan tensor. For more general tensors, a decomposition of the whole TN in three-legged terms can be used [107], or more efficient precomputation of the corresponding coefficients can be done in the algorithm [111, 113, 114]. Notice that generic tensors (i.e. without explicit symmetry) can also be used to describe a TNS with the desired global symmetry, even producing a more compact description [115]. Using the symmetry structure of the tensors involves a more cumbersome implementation of the methods (described in detail in the previous references), but in exchange allows one to work with blocks which have smaller bond dimension, which reduces the computational cost of contractions at the lowest-level.

Symmetric tensors can be used to raise the global symmetry to a gauge one [116–118]. This is done through the introduction of additional link tensors (analogous to link variables in usual formulations of lattice gauge theories).

Finally, it is worth mentioning that, at the theoretical level, a framework has been developed to characterize MPS and PEPS in terms of the tensor symmetries [37], a formal approach that has produced fundamental results and continues to be an active and fruitful area of research [24].

C. Fermions

An advantage of the TN framework with respect to other numerical methods for quantum many-body problems, is the possibility to treat problems with fermionic degrees of freedom, of fundamental interest for condensed matter and fundamental physics. Whereas in this case Quantum Monte Carlo methods are often obstructed by the sign problem, which causes the cost of convergence to increase exponentially with the system size, TN calculations can indistinctly treat fermionic and spin setups.

In one spatial dimension, fermionic modes do not pose a real problem, as they can be mapped to spins through the Jordan-Wigner transformation. This maps local fermionic models onto local spin Hamiltonians, such that both can be treated with exactly the same algorithms. In higher dimensions, however, a similar transformation does not maintain the locality of the model. An alternative that maps local fermions to local spins and would support a treatment with standard TNS algorithms was introduced in [119], but at the cost of introducing additional degrees of freedom doubling the size of the system.

It is however possible to define TNS directly in terms of fermionic degrees of freedom. The explicit construction was presented by several independent, but essentially equivalent, proposals [83, 103, 120, 121]. The fundamental idea is to work in a representation in which all spaces, virtual and physical, are fermionic, and have well-defined parity, i.e. the tensors are symmetric with respect to parity transformations, in the sense described in III B. Then it is possible to encode the statistics of fermionic operators in a local way, such that the scaling of the computational cost with the system size is preserved.

The most intuitive formulation [83, 120] can be visualized as an effective linear ordering of the fermionic modes, fixed once a graphical representation of the TNS is chosen (the order would be that obtained when projecting all the sites of the graph onto a line). Each crossing of legs in the diagram has to be accounted for, as it involves commutations of fermionic operators. This can be achieved substituting the crossing by a swap matrix, which introduces a negative sign with fermionic degrees of freedom with odd parity are exchanged. Thanks to the symmetry of the tensors, the swap matrices can be moved through the network and be absorbed into local tensors, and the contraction can follow the same sequence as in the spin case, thus keeping the leading cost. This formalism, which can be combined with additional symmetries [114], has already made possible for iPEPS to

beat any other computational method in some parameter regimes of the Hubbard model [94].

D. Dynamics

Simulating time evolution is a crucial tool for understanding the out-of-equilibrium dynamics of quantum many-body systems, linked to fundamental questions such as thermalization. Together with the applicability to fermionic problems, being able to address real-time evolution is precisely one of the main advantages of TNS as compared to Monte Carlo methods.

The TNS toolbox has several different methods to tackle these problems. Many of them produce an approximation to the time-evolved state within the desired family [see [72] for a recent detailed review]. The standard algorithms described in IIC 2 proceed by constructing an approximation of the evolution operator $U(\delta) = e^{-i\delta H}$ for a finite time step δ and applying it onto a TNS wave function. In general, this increases the bond dimension, and it must be followed by a truncation step that reduces the tensors again. A limitation of these methods is that they rely on approximations of the Hamiltonian exponential operator that become exceedingly costly as the range of the interactions increases.

Krylov-based methods, instead, directly target the result of the evolution step by approximating the application of the operator on the state as a linear combination of Krylov vectors [63, 122], instead of explicitly approximating the evolution operator in the full space. This in turn requires approximating the Krylov vectors themselves by TNS. A related approach is using Chebyshev expansions of the exponential operators [123].

Also the more recently proposed time dependent variational principle (TDVP) [124, 125] adopts a different strategy, in which the MPS tensors are evolved such that the evolution never leaves the MPS manifold. This is achieved by projecting the variation of the wave function, given by the rhs of the Schrödinger equation, onto the local tangent plane of the MPS. Despite its different philosophy, TDVP algorithms for finite and infinite systems can be formulated in terms of essentially the same low level primitives as the traditional ones [125]. That is, the tensors of the ansatz can be updated according to the solution of a local evolution, in this case given by effective Hamiltonians that result from the tangent plane projection. An advantage of this method is that it preserves conserved quantities of the evolved state, such as the norm and energy.

In the uniform MPS case, the TDVP algorithm is the first exponent of a new generation of TNS algorithms, so-called tangent-space methods [126], based on exploiting the geometric structure of the MPS manifold. These increasingly popular methods have multiple applications beyond time evolution, including the variational optimization of uMPS or finding elementary excitations, and have been partly adapted for PEPS [see [127] for a ped-

agogical overview]. More recently, generalizations for TTN and other isometric TN have been introduced [128–130].

The approaches described above provide powerful algorithms to investigate the evolution of quantum systems for moderate times, or close to equilibrium [72]. However, they are still subject to the fundamental limitation mentioned in IIC 2: under time evolution, entanglement can grow fast, the bond dimension of the ansatz would need to grow exponentially with the simulated time [70, 71], such that after short times, the simulation becomes unfeasible, a problem that has been termed *entanglement barrier*. But for physical problems, often the interest is not in the full description of the state, but in expectation values of local observables, which correspond to experimentally accessible quantities. There a paradoxical situation takes place, since in the long time limit observables are expected to thermalize or equilibrate to values that are well described by statistical ensembles, which can be themselves efficiently approximated by (mixed) TNS, but in most cases the entanglement barrier makes it impossible to reach this regime following the evolution of the state.

For this reason, an active effort is being dedicated to the investigation of potentially new methods that avoid the entanglement barrier and manage to describe the long time dynamics of local quantities. A first proposal was evolving operators in Heisenberg picture [131] using a suitably adapted time evolution algorithm. Despite not completely solving the entanglement problem, such an approach constitutes the basis of many other strategies for dynamical quantities. Another idea was to target the TN that represents the time-dependent local observables, and to approximate its contraction in the transverse direction, after folding [132, 133], which can give access to longer times, especially when exploiting the finite propagation velocity of correlations [134, 135]. This remains an active area of research, and several new strategies have been proposed in the last years to focus on the local observables [136–138].

E. Excitations

With the variational approach for the ground state (section IIC 1) it is possible to target also low excited states, by simply orthogonalizing the targeted state with respect to any number of previously computed ones [16, 26], an approach that is most useful in the case of finite systems.

A particularly useful ansatz for elementary excitations is to model them as local perturbations acting on the vacuum. In the TNS framework, it is possible to construct well-defined momentum states of this form by suitable superpositions of a locally modified ground state [7]. Tangent-space methods offer a way to generalize this construction that is especially powerful in the thermodynamic limit [126, 127]. In this framework, elementary

excitations are written as tangent vectors with position-dependent momentum factors, and their energies can be optimized variationally. Also topologically non-trivial excitations (such as domain walls) can be captured in this language.

Although low-energy excitations as the ones above are often observed to fulfill an approximate area law, the same is not true for generic, highly excited states. An exception is the case of many-body localized Hamiltonians. Hence, several specific algorithms have been developed to target eigenstates at some high energy value E , for instance using a *shift and invert* strategy [139], targeting the state at given energy that maximizes the overlap with a particular product state [140], or searching for the lowest eigenvalue of $(H - E)^2$ [141].

IV. FURTHER TN APPROACHES AND PERSPECTIVES

Other aspects of TN technologies, beyond the standard TNS tools discussed in the previous sections offer additional ways to explore the physics of complex systems.

A. Network renormalization approaches

Some of the earliest works in the TN literature, before the quantum information perspective shaped the language for TN, already pointed out the connection between many-body problems and tensors in the partition functions of classical spin systems [10, 142]. In this approach, a TN represents exactly the partition function of a classical model (which might as well correspond to a path integral formulation of a quantum one) and tensor contractions can be used to approximate a result.

The tensor renormalization group (TRG) method introduced in [143] is based on a block renormalization of a two-dimensional TN: in each coarse-graining step a local group of tensors is replaced by their approximate contraction with truncated bonds, such that the size of the TN is divided by a constant (see fig. 4). The original truncation is done by simply singular value decompositions of the tensors being contracted. In [144] a new strategy was introduced, called second renormalization group (SRG) method, where a different truncation is chosen that tries to maintain the fidelity of the contraction of the whole network, by taking into account the environment of the tensor that is being computed. A more efficient contraction and truncation strategy that can be applied to higher dimensional systems was later proposed in [145], using the higher order singular value decomposition.

A shortcoming of the approach, already identified in [143] is that some short-range entanglement structures cannot be removed by the TRG coarse-graining,

in particular, the corner double line (CDL) tensor. Several modifications have been proposed to solve this issue, such as the TNR (tensor network renormalization) that includes disentangles, in the spirit of MERA, before the renormalization steps [146]. Other proposals have been the iterative optimization of the tensors around a loop [147], or different local index truncations that take care of internal correlations [86, 148].

TRG approaches are also useful to contract the TN corresponding to observables for quantum states in higher dimensions, and can then be used as part of PEPS optimization algorithms [80, 81] (see sec. III A). A related topic is the treatment of fermionic problems in TRG approaches. In [149] it was shown that wave functions and expectation values of many-body fermionic (but also bosonic) systems could be expressed and contracted as a Grassmann tensor network, in which tensor components are given in terms of Grassmann variables, and for which a suitable TRG approach can be defined. A compact ansatz of this form, together with algorithms to renormalize the network and to evolve the tensors were presented in [150], and have been used, for instance, to study discretized field theories with fermionic degrees of freedom [see references in [151, 152]].

B. Connections to other techniques

Exploring the potential connections between TN methods and other techniques is an exciting possibility that, on the one hand, can result in new or improved algorithms and, on the other, opens the door to treating new problems with TN methods, as the following examples illustrate.

- Monte Carlo algorithms. Monte Carlo sampling can be used to speed up TN contractions, and variationally optimize TNS parameters [153, 154]. With a complementary perspective, TN contractions can be employed to directly sample configurations from the partition function [155–157], but also to define a Markov chain with collective updates [158].
- Machine Learning. The connections between TN and machine learning drive some of the most recent developments, including the use of TNS models for machine learning tasks [159, 160], and also importing numerical tools, such as automatic differentiation, into TN algorithms [161].
- Field theory. The interplay between TNS and quantum field theory is another decidedly active area, which has produced accurate numerical results for lattice gauge theories [151, 152], but also motivates formal developments, such as gauge symmetric (see sec. III B) and continuous [162, 163] formulations of TNS.




FIG. 4. Coarse-graining step in the simplest TRG schemes: (a) TN representing a partition function of a classical spin model; (b) original TRG; (c) higher order TRG (HOTRG).

V. OUTLOOK

The field of tensor networks has grown impressively in the last decade and remains a vibrant research area. Current TN research moves forward in different directions. A rather formal approach explores the mathematical aspects of these ansatzes. With a more applied perspective, significant effort is being devoted to the development of numerical TN methods, a multifaceted enterprise, some of whose spotlights have been highlighted in the previous pages. And the field continues to uncover synergies with seemingly remote topics, and to develop in new and creative ways. All these directions are likely to produce exciting results in the coming years, maybe finding useful TNS subfamilies, improving the efficiency of high-dimensional or dynamical calculations, or bridging the gaps between formal and numerical developments.

At the same time, mature TN algorithms are well established as competitive computational methods for the study of many-body problems. These algorithms, reviewed in the first part of this article, make it easy for the newcomer to try TN for an existing problem, and simultaneously serve as a platform for the more specialized researcher to experiment with new algorithms or to draw new connections between TN and other disciplines.

ACKNOWLEDGMENTS

I am deeply grateful to E. Carmona, P. Emonts, M. Frías-Pérez and T. Nishino for their critical reading and constructive comments on an earlier version of this article. This work was partly funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2111 – 390814868.

-
- [1] Baxter RJ. 1968. *J. Math. Phys.* 9(4):650–654
 - [2] Affleck I, Kennedy T, Lieb EH, Tasaki H. 1987. *Phys. Rev. Lett.* 59(7):799–802
 - [3] Okunishi K, Nishino T, Ueda H. 2021
 - [4] Klümper A, Schadschneider A, Zittartz J. 1993. *Europhysics Letters (EPL)* 24(4):293–297
 - [5] Fannes M, Nachtergaele B, Werner RF. 1992. *Commun. Math. Phys.* 144(3):443–490
 - [6] White SR. 1992. *Phys. Rev. Lett.* 69(19):2863–2866
 - [7] Östlund S, Rommer S. 1995. *Phys. Rev. Lett.* 75(19):3537–3540
 - [8] Dukelsky J, Martín-Delgado MA, Nishino T, Sierra G. 1998. *Europhysics Letters (EPL)* 43(4):457–462
 - [9] Nishino T, Okunishi K. 1995. *J. Phys. Soc. Jpn.* 64(11):4084–4087
 - [10] Nishino T, Okunishi K. 1996. *J. Phys. Soc. Jpn.* 65(4):891–894
 - [11] Nishino T, Okunishi K. 1998. *Journal of the Physical Society of Japan* 67(9):3066–3072
 - [12] Hallberg K. 2003. *arXiv e-prints :cond-mat/0303557*
 - [13] Schollwöck U. 2005. *Rev. Mod. Phys.* 77(1):259–315
 - [14] Vidal G. 2003. *Phys. Rev. Lett.* 91(14):147902
 - [15] Verstraete F, Porras D, Cirac JI. 2004. *Phys. Rev. Lett.* 93(22):227205
 - [16] McCulloch IP. 2007. *J. Stat. Mech.* 2007(10):P10014
 - [17] Vidal G. 2004. *Phys. Rev. Lett.* 93(4):040502
 - [18] Daley AJ, Kollath C, Schollwöck U, Vidal G. 2004. *J. Stat. Mech.* 2004(04):P04005
 - [19] White SR, Feiguin AE. 2004. *Phys. Rev. Lett.* 93(7):076401
 - [20] Verstraete F, García-Ripoll JJ, Cirac JI. 2004. *Phys. Rev. Lett.* 93(20):207204
 - [21] Zwolak M, Vidal G. 2004. *Phys. Rev. Lett.* 93(20):207205
 - [22] Feiguin AE, White SR. 2005. *Phys. Rev. B* 72(22):220401
 - [23] Verstraete F, Cirac JI. 2004. *arXiv:cond-mat/0407066*
 - [24] Cirac JI, Pérez-García D, Schuch N, Verstraete F. 2021. *Rev. Mod. Phys.* 93(4):045003
 - [25] Verstraete F, Murg V, Cirac J. 2008. *Adv. Phys.* 57(2):143–224
 - [26] Schollwöck U. 2011. *Ann. Phys.* 326(1):96 – 192 January 2011 Special Issue
 - [27] Orús R. 2014. *Annals Phys.* 349:117–158
 - [28] Bridgeman JC, Chubb CT. 2017. *J. Phys. A: Math. Theor.* 50(22):223001
 - [29] Ran SJ, Tirrito E, Peng C, Chen X, Tagliacozzo L, et al. 2020. Tensor network contractions. Lecture Notes

- in Physics. Springer International Publishing
- [30] Silvi P, Tschirsich F, Gerster M, Jünemann J, Jaschke D, et al. 2019. *SciPost Phys. Lect. Notes* :8
- [31] Eisert J, Cramer M, Plenio MB. 2010. *Rev. Mod. Phys.* 82(1):277–306
- [32] Hastings MB. 2007. *J. Stat. Mech.* 2007(08):P08024
- [33] Wolf MM, Verstraete F, Hastings MB, Cirac JI. 2008. *Phys. Rev. Lett.* 100(7):070502
- [34] Calabrese P, Cardy J. 2004. *J. Stat. Mech.* 2004(06):P06002
- [35] Wolf MM. 2006. *Phys. Rev. Lett.* 96(1):010404
- [36] Pérez-García D, Verstraete F, Wolf MM, Cirac JI. 2007. *Quantum Inf. Comput.* 7:401–430
- [37] Schuch N, Cirac I, Pérez-García D. 2010. *Annals Phys.* 325(10):2153 – 2192
- [38] Schuch N, Wolf MM, Verstraete F, Cirac JI. 2007. *Phys. Rev. Lett.* 98(14):140506
- [39] Verstraete F, Wolf MM, Pérez-García D, Cirac JI. 2006. *Phys. Rev. Lett.* 96(22):220601
- [40] Shi YY, Duan LM, Vidal G. 2006. *Phys. Rev. A* 74(2):022320
- [41] Cirac JI, Verstraete F. 2009. *J. Phys. A: Math. Theor.* 42(50):504004
- [42] Silvi P, Giovannetti V, Montangero S, Rizzi M, Cirac JI, Fazio R. 2010. *Phys. Rev. A* 81(6):062335
- [43] Tagliacozzo L, Evenbly G, Vidal G. 2009. *Phys. Rev. B* 80(23):235127
- [44] Vidal G. 2007. *Phys. Rev. Lett.* 99(22):220405
- [45] Vidal G. 2008. *Phys. Rev. Lett.* 101(11):110501
- [46] Evenbly G, Vidal G. 2009. *Phys. Rev. B* 79(14):144108
- [47] Pfeifer RNC, Evenbly G, Vidal G. 2009. *Phys. Rev. A* 79(4):040301
- [48] Montangero S, Rizzi M, Giovannetti V, Fazio R. 2009. *Phys. Rev. B* 80(11):113103
- [49] Barthel T, Kliesch M, Eisert J. 2010. *Phys. Rev. Lett.* 105(1):010502
- [50] Evenbly G, Vidal G. 2014. *Phys. Rev. Lett.* 112(24):240502
- [51] Pirvu B, Murg V, Cirac JI, Verstraete F. 2010. *New J. Phys.* 12(2):025012
- [52] de las Cuevas G, Schuch N, Pérez-García D, Cirac JI. 2013. *New J. Phys.* 15(12):123021
- [53] Pfeifer RNC, Haegeman J, Verstraete F. 2014. *Phys. Rev. E* 90(3):033315
- [54] Gray J, Kourtis S. 2021. *Quantum* 5:410
- [55] Note1. ????. This corresponds to the single site algorithm, most natural in the TN framework. Some modifications can be made to connect to the classic two-site DMRG [see [3, 26] for details on DMRG variants and their historical development].
- [56] Hubig C, McCulloch IP, Schollwöck U. 2017. *Phys. Rev. B* 95(3):035129
- [57] Pollmann F, Mukerjee S, Turner AM, Moore JE. 2009. *Phys. Rev. Lett.* 102(25):255701
- [58] Pirvu B, Vidal G, Verstraete F, Tagliacozzo L. 2012. *Phys. Rev. B* 86(7):075117
- [59] Cui J, Cirac JI, Bañuls MC. 2015. *Phys. Rev. Lett.* 114(22):220601
- [60] Note2. ????. Notice that the term is used loosely in the literature, sometimes interchanged with tDMRG—see [26] for the details.
- [61] Zaletel MP, Mong RSK, Karrasch C, Moore JE, Pollmann F. 2015. *Phys. Rev. B* 91(16):165112
- [62] White SR. 2009. *Phys. Rev. Lett.* 102(19):190601
- [63] García-Ripoll JJ. 2006. *New J. Phys.* 8(12):305
- [64] Huckle T, Waldherr K. 2012. *PAMM* 12(1):641–642
- [65] Vidal G. 2007. *Phys. Rev. Lett.* 98(7):070201
- [66] Dubail J. 2017. *J. Phys. A: Math. Theor.* 50(23):234001
- [67] Hastings MB. 2006. *Phys. Rev. B* 73(8):085115
- [68] Molnar A, Schuch N, Verstraete F, Cirac JI. 2015. *Phys. Rev. B* 91(4):045138
- [69] Calabrese P, Cardy J. 2005. *J. Stat. Mech.: Theory Exp.* 2005(04):P04010
- [70] Osborne TJ. 2006. *Phys. Rev. Lett.* 97(15):157202
- [71] Schuch N, Wolf MM, Verstraete F, Cirac JI. 2008. *Phys. Rev. Lett.* 100(3):030504
- [72] Paeckel S, Köhler T, Swoboda A, Manmana SR, Schollwöck U, Hubig C. 2019. *Annals Phys.* 411:167998
- [73] Stoudenmire E, White SR. 2012. *Annu. Rev. Condens. Matter Phys.* 3(1):111–128
- [74] Cataldi G, Abedi A, Magnifico G, Notarnicola S, Pozza ND, et al. 2021. *Quantum* 5:556
- [75] Yan S, Huse DA, White SR. 2011. *Science* 332(6034):1173–1176
- [76] Depenbrock S, McCulloch IP, Schollwöck U. 2012. *Phys. Rev. Lett.* 109(6):067201
- [77] Murg V, Verstraete F, Cirac JI. 2007. *Phys. Rev. A* 75(3):033605
- [78] Jordan J, Orús R, Vidal G, Verstraete F, Cirac JI. 2008. *Phys. Rev. Lett.* 101(25):250602
- [79] Lubasch M, Cirac JI, Bañuls MC. 2014. *New J. Phys.* 16(3):033014
- [80] Jiang HC, Weng ZY, Xiang T. 2008. *Phys. Rev. Lett.* 101(9):090603
- [81] Gu ZC, Levin M, Wen XG. 2008. *Phys. Rev. B* 78(20):205116
- [82] Orús R, Vidal G. 2009. *Phys. Rev. B* 80(9):094403
- [83] Corboz P, Orús R, Bauer B, Vidal G. 2010. *Phys. Rev. B* 81(16):165104
- [84] Lubasch M, Cirac JI, Bañuls MC. 2014. *Phys. Rev. B* 90(6):064425
- [85] Phien HN, Bengua JA, Tuan HD, Corboz P, Orús R. 2015. *Phys. Rev. B* 92(3):035142
- [86] Evenbly G. 2018. *Phys. Rev. B* 98(8):085155
- [87] Osorio Iregui J, Troyer M, Corboz P. 2017. *Phys. Rev. B* 96(11):115113
- [88] Note3. ????. Notice that, although finite-size extrapolation from PEPS is possible, the number of tensors to determine (e.g. L^2 for a two-dimensional system) makes the calculations exceedingly long already for relatively small sizes.
- [89] Corboz P. 2016. *Phys. Rev. B* 94(3):035133
- [90] Vanderstraeten L, Haegeman J, Corboz P, Verstraete F. 2016. *Phys. Rev. B* 94(15):155123
- [91] Rader M, Läuchli AM. 2018. *Phys. Rev. X* 8(3):031030
- [92] Corboz P, Czarnik P, Kapteijns G, Tagliacozzo L. 2018. *Phys. Rev. X* 8(3):031031
- [93] Vanhecke B, Hasik J, Verstraete F, Vanderstraeten L. 2021. *arXiv e-prints* :arXiv:2102.03143
- [94] Zheng BX, Chung CM, Corboz P, Ehlers G, Qin MP, et al. 2017. *Science* 358(6367):1155–1160
- [95] Vlaar PCG, Corboz P. 2021. *Phys. Rev. B* 103(20):205137
- [96] Czarnik P, Dziarmaga J, Corboz P. 2019. *Phys. Rev. B* 99(3):035115
- [97] Kshetrimayum A, Weimer H, Orús R. 2017. *Nat. Commun.* 8(1):1291

- [98] Kilda D, Biella A, Schiro M, Fazio R, Keeling J. 2021. *SciPost Phys. Core* 4(1):5
- [99] Mc Keever C, Szymańska MH. 2021. *Phys. Rev. X* 11(2):021035
- [100] Hubig C, Cirac JI. 2019. *SciPost Phys.* 6(3):31
- [101] Bañuls MC, Pérez-García D, Wolf MM, Verstraete F, Cirac JI. 2008. *Phys. Rev. A* 77(5):052306
- [102] Zaletel MP, Pollmann F. 2020. *Phys. Rev. Lett.* 124(3):037201
- [103] Kraus CV, Schuch N, Verstraete F, Cirac JI. 2010. *Phys. Rev. A* 81(5):052338
- [104] Magnifico G, Felser T, Silvi P, Montangero S. 2021. *Nat. Commun.* 12(1):3600
- [105] Felser T, Notarnicola S, Montangero S. 2021. *Phys. Rev. Lett.* 126(17):170603
- [106] McCulloch IP, Gulácsi M. 2002. *Europhysics Letters (EPL)* 57(6):852–858
- [107] Singh S, Pfeifer RNC, Vidal G. 2010. *Phys. Rev. A* 82(5):050301
- [108] Singh S, Pfeifer RNC, Vidal G. 2011. *Phys. Rev. B* 83(11):115125
- [109] Bauer B, Corboz P, Orús R, Troyer M. 2011. *Phys. Rev. B* 83(12):125106
- [110] Singh S, Vidal G. 2012. *Phys. Rev. B* 86(19):195114
- [111] Weichselbaum A. 2012. *Ann. Phys.* 327(12):2972 – 3047
- [112] Note4. ????. In the non-Abelian case, q will actually be a composite index, including not only the label for the irrep, but also additional quantum numbers to account for its inner (and potentially outer) multiplicity [114].
- [113] Hubig C. 2018. *SciPost Phys.* 5(5):47
- [114] Bruognolo B, Li JW, von Delft J, Weichselbaum A. 2021. *SciPost Phys. Lect. Notes* :25
- [115] Singh S, Vidal G. 2013. *Phys. Rev. B* 88(11):115147
- [116] Tagliacozzo L, Celi A, Lewenstein M. 2014. *Phys. Rev. X* 4(4):041024
- [117] Zohar E, Burrello M. 2016. *New J. Phys.* 18(4):043008
- [118] Haegeman J, Van Acoleyen K, Schuch N, Cirac JI, Verstraete F. 2015. *Phys. Rev. X* 5(1):011024
- [119] Verstraete F, Cirac JI. 2005. *J. Stat. Mech.: Theory Exp.* 2005(09):P09012–P09012
- [120] Corboz P, Vidal G. 2009. *Phys. Rev. B* 80(16):165129
- [121] Pineda C, Barthel T, Eisert J. 2010. *Phys. Rev. A* 81(5):050303
- [122] Wall ML, Carr LD. 2012. *New J. Phys.* 14(12):125015
- [123] Holzner A, Weichselbaum A, McCulloch IP, Schollwöck U, von Delft J. 2011. *Phys. Rev. B* 83(19):195115
- [124] Haegeman J, Cirac JI, Osborne TJ, Pižorn I, Verschelde H, Verstraete F. 2011. *Phys. Rev. Lett.* 107(7):070601
- [125] Haegeman J, Lubich C, Oseledets I, Vandereycken B, Verstraete F. 2016. *Phys. Rev. B* 94(16):165116
- [126] Haegeman J, Osborne TJ, Verstraete F. 2013. *Phys. Rev. B* 88(7):075133
- [127] Vanderstraeten L, Haegeman J, Verstraete F. 2019. *SciPost Phys. Lect. Notes* :7
- [128] Kloss B, Reichman DR, Lev YB. 2020. *SciPost Phys.* 9(5):70
- [129] Bauernfeind D, Aichhorn M. 2020. *SciPost Phys.* 8(2):24
- [130] Hauru M, Damme MV, Haegeman J. 2021. *SciPost Phys.* 10(2):40
- [131] Hartmann MJ, Prior J, Clark SR, Plenio MB. 2009. *Phys. Rev. Lett.* 102(5):057202
- [132] Bañuls MC, Hastings MB, Verstraete F, Cirac JI. 2009. *Phys. Rev. Lett.* 102(24):240603
- [133] Müller-Hermes A, Cirac JI, Bañuls MC. 2012. *New J. Phys.* 14(7):075003
- [134] Frias-Pérez M, Bañuls MC. 2022
- [135] Leroise A, Sonner M, Abanin DA. 2022
- [136] White CD, Zaletel M, Mong RSK, Refael G. 2018. *Phys. Rev. B* 97(3):035127
- [137] Surace J, Piani M, Tagliacozzo L. 2019. *Phys. Rev. B* 99(23):235115
- [138] Rakovszky T, von Keyserlingk CW, Pollmann F. 2022. *Phys. Rev. B* 105(7):075131
- [139] Yu X, Pekker D, Clark BK. 2017. *Phys. Rev. Lett.* 118(1):017201
- [140] Khemani V, Pollmann F, Sondhi SL. 2016. *Phys. Rev. Lett.* 116(24):247204
- [141] Lim SP, Sheng DN. 2016. *Phys. Rev. B* 94(4):045111
- [142] Nishino T, Hieida Y, Okunishi K, Maeshima N, Akutsu Y, Gendiar A. 2001. *Progress of Theoretical Physics* 105(3):409–417
- [143] Levin M, Nave CP. 2007. *Phys. Rev. Lett.* 99(12):120601
- [144] Xie ZY, Jiang HC, Chen QN, Weng ZY, Xiang T. 2009. *Phys. Rev. Lett.* 103(16):160601
- [145] Xie ZY, Chen J, Qin MP, Zhu JW, Yang LP, Xiang T. 2012. *Phys. Rev. B*
- [146] Evenbly G, Vidal G. 2015. *Phys. Rev. Lett.* 115(18):180405
- [147] Yang S, Gu ZC, Wen XG. 2015. *arXiv:1512.04938*
- [148] Hauru M, Delcamp C, Mizera S. 2018. *Phys. Rev. B* 97(4):045111
- [149] Gu ZC, Verstraete F, Wen XG. 2010
- [150] Gu ZC. 2013. *Phys. Rev. B* 88(11):115139
- [151] Bañuls MC, Cichy K. 2020. *Reports on Progress in Physics* 83(2):024401
- [152] Meurice Y, Sakai R, Unmuth-Yockey J. 2020. *Rev. Mod. Phys.*
- [153] Sandvik AW, Vidal G. 2007. *Phys. Rev. Lett.* 99(22):220602
- [154] Wang L, Pižorn I, Verstraete F. 2011. *Phys. Rev. B* 83(13):134421
- [155] Ueda K, Otani R, Nishio Y, Gendiar A, Nishino T. 2005. *J. Phys. Soc. Jpn.* 74(Suppl):111–114
- [156] Ferris AJ, Vidal G. 2012. *Phys. Rev. B* 85(16):165146
- [157] Rams MM, Mohseni M, Eppens D, Jałowiecki K, Gardas B. 2021. *Phys. Rev. E* 104(2):025308
- [158] Frías-Pérez M, Mariën M, Pérez-García D, Bañuls MC, Iblisdir S. 2021
- [159] Stoudenmire E, Schwab DJ. 2016. In *Advances in Neural Information Processing Systems*, eds. D Lee, M Sugiyama, U Luxburg, I Guyon, R Garnett, vol. 29. Curran Associates, Inc.
- [160] Han ZY, Wang J, Fan H, Wang L, Zhang P. 2018. *Phys. Rev. X* 8(3):031012
- [161] Liao HJ, Liu JG, Wang L, Xiang T. 2019. *Phys. Rev. X* 9(3):031041
- [162] Verstraete F, Cirac JI. 2010. *Phys. Rev. Lett.* 104(19):190405
- [163] Haegeman J, Osborne TJ, Verschelde H, Verstraete F. 2013. *Phys. Rev. Lett.* 110(10):100402

RELATED RESOURCES

1. T. Nishino, The DMRG Homepage, <http://quattro.phys.sci.kobe-u.ac.jp/dmrg.html>

2. E. M. Stoudenmire, The Tensor Network, <https://tensornetwork.org/>
3. M. Fishman, E. M. Stoudenmire, S. R.White, An open source TNS library, <https://itensor.org/>
4. J. Hauschild, F. Pollmann, A tensor library for Python (TeNPy), <https://github.com/tenpy/tenpy>
5. G. Evenbly, Tensors.net, <https://www.tensors.net/>