Faster identification of optimal contraction sequences for tensor networks

Robert N. C. Pfeifer,^{1,*} Jutho Haegeman,² and Frank Verstraete^{2,3}

¹*Perimeter Institute for Theoretical Physics, 31 Caroline St. N, Waterloo Ontario N2L 2Y5, Canada*

²*Department of Physics and Astronomy, Ghent University, Krijgslaan 281-S9, B-9000 Ghent, Belgium*

³*Vienna Center for Quantum Science, Universität Wien, Boltzmanngasse 5, A-1090 Wien, Austria*

(Received 22 April 2013; revised manuscript received 1 July 2014; published 30 September 2014)

The efficient evaluation of tensor expressions involving sums over multiple indices is of significant importance to many fields of research, including quantum many-body physics, loop quantum gravity, and quantum chemistry. The computational cost of evaluating an expression may depend strongly on the order in which the index sums are evaluated, and determination of the operation-minimizing contraction sequence for a single tensor network (single term, in quantum chemistry) is known to be NP-hard. The current preferred solution is an exhaustive search, using either an iterative depth-first approach with pruning or dynamic programming and memoization, but these approaches are impractical for many of the larger tensor network ansätze encountered in quantum many-body physics. We present a modified search algorithm with enhanced pruning which exhibits a performance increase of several orders of magnitude while still guaranteeing identification of an optimal operation-minimizing contraction sequence for a single tensor network. A reference implementation for MATLAB, compatible with the `ncon()` and `multienv()` network contractors of [arXiv:1402.0939](https://arxiv.org/abs/1402.0939) and Evenbly and Pfeifer, *Phys. Rev. B* **89**, 245118 (2014), respectively, is supplied.

DOI: [10.1103/PhysRevE.90.033315](https://doi.org/10.1103/PhysRevE.90.033315)

PACS number(s): 02.70.-c, 05.30.-d, 03.67.-a

I. INTRODUCTION

A. Overview

The need to efficiently contract a tensor expression is one which arises in many different areas of research, including both quantum chemistry and quantum many-body physics, and has been an area of intense study since at least 1997 (see, e.g., Refs. [1–17]), despite being known to be NP-hard [14]. In quantum many-body physics interest in this problem has been driven by the increasing complexity of tensor network ansätze for many-body systems (see Sec. IB), while in quantum physical chemistry it has long been acknowledged that determination of the optimal contraction sequences for increasingly complex tensor networks represents a significant bottleneck in the development of new algorithms [1,2,4,13].

In its most general form, the problem is to evaluate a multidimensional tensor sum such as

$$\sum_{j,k,m,p} A_{ijk} B_{jlm} C_{kmp} D_{pq} + \sum_{j,k} A_{ijk} D_{jklq} \quad (1)$$

as rapidly as possible subject to the constraints of available computing hardware. This problem may be seen as a generalization of the matrix-chain multiplication problem, described in Appendix A, where a string of matrices are to be multiplied together as efficiently as possible. Unlike the matrix-chain multiplication problem, however, this problem cannot be solved in polynomial time through the use of dynamic programming techniques [14].

While this optimization problem is intrinsically multidimensional, balancing available memory and (for multinode machines) internode communication delays against the number of floating-point operations which must be performed, the predominant approach to this problem is, first, to identify the ideal contraction procedure which would be performed

on a single node with infinite resources, minimizing the number of floating point operations to be performed (*operation minimization*), before trading off performance against memory constraints and distributing the problem across multiple nodes [2,7,8,12,15,17]. Consequently, the task of operation minimization is of fundamental importance.

A significant difference between applications in quantum chemistry and quantum many-body physics is that the ansätze of the former frequently yield tensor expressions involving many summed terms which may be factorized in numerous different ways, and it is necessary to explore both the different factorizations and the different index contraction sequences for a given factorization, with the latter task having been termed *single term optimization* [10]. On the other hand, in quantum many-body physics an ansatz is typically made up of a single term, but this term may involve substantially more factors than are encountered in the single terms of quantum chemistry. The problem of factorizing multiple-term tensor expressions for optimal computational efficiency has been explored in depth elsewhere [10,11] and itself depends on the determination of efficient contraction sequences for single terms, so the task of single term operation minimization is of key importance in both quantum chemistry and quantum many-body physics (where it is synonymous with *optimal contraction of a tensor network*). It is the problem of rapid single term optimization which is addressed in this paper.

As the primary interest of the authors is in quantum many-body physics, we will frequently favor quantum physics terminology over that of quantum chemistry or linear algebra. It should therefore be noted that a *tensor network* is synonymous with a *single term in a tensor expression* and that the *optimal* contraction sequence is assumed to be that which minimizes the number of floating point operations performed. For a single term, if it is assumed that evaluation of

$$C_{ij} = \sum_k A_{ik} \times B_{kj} \quad (2)$$

*robert.pfeifer@mq.edu.au

begins by constructing an array of appropriate size for matrix C_{ij} and populated with zeros, then minimization of the number of floating point operations is synonymous with minimizing the number of floating point multiplications as the ratio of multiplications to additions is 1:1. We note that identifying an optimal contraction sequence for a single tensor network corresponds to solving the “multiplication problem” of Ref. [14] and thus is known to be NP-hard.

An index i runs from 1 to $|i|$, where $|i|$ is termed the *dimension* of index i and may also be denoted ξ_i . The dimension of a tensor, e.g., $|C_{ijk}|$, is then the product of the dimension of its indices and corresponds to the number of entries in the multidimensional array C_{ijk} .

For complicated tensor networks involving large numbers of tensors and index sums, we will use a graphical notation which is a simplification of that of Penrose [18] and is summarized in Sec. 1.2 of Ref. [19]. In this notation, shapes represent tensors and lines (or *legs*) represent indices. A line connecting two shapes therefore represents an index appearing on the tensors corresponding to both shapes. We follow the Einstein summation convention, where any repeated index appearing once in the upper position and once in the lower position is assumed summed, unless otherwise specified. The multiplication of two matrices $C = A \times B$ may therefore be written as

$$C_{ij} = A_{ik} B_{kj} \quad (3)$$

and represented graphically as in Fig. 1(a), while the inner product of a matrix and two column vectors, $v^T M w$, may be written as

$$c = v_i M^{ij} w_j \quad (4)$$

and represented graphically as in Fig. 1(b).

B. Tensor network algorithms in quantum many-body physics

Tensor network algorithms provide powerful tools for the study of a wide variety of physical systems. They are perhaps best known for their use in condensed-matter physics as numerical techniques for the study of quantum many-body systems on a lattice (e.g., Refs. [20–44]), but recent breakthroughs blending ideas from quantum information with advanced numerical techniques have led to the construction of new ansätze (e.g., Refs. [26,27,45–48]) having applications in fields as diverse as holography and the anti-de Sitter–conformal field theory correspondence [49–51], the study of many-body entanglement [33,45,48,52], and the classification of topological phases in quantum spin systems [53–56]. As a numerical tool, a tensor network algorithm typically comprises an ansatz for the description of pure or mixed quantum states, which is composed of a network of tensors, and an iterative

procedure for updating this ansatz. Examples include the density matrix renormalization group (DMRG) [57,58] and time evolving block decimation (TEBD) algorithms [59,60], both of which are based on the matrix product state (MPS) ansatz, and tree tensor networks (TTNs) [61], projected entangled pair states (PEPS) [62,63], and the multiscale entanglement renormalization ansatz (MERA) [26,27,45–48,64].

The fundamental challenge to the numerical study of quantum many-body systems on a lattice is that the number of degrees of freedom, and thus the computational cost associated with exact simulation, grows exponentially with the size of the system. To overcome this challenge, tensor network ansätze replace the coefficients $c_{i_1 \dots i_n}$ of a quantum state $|\psi\rangle$,

$$|\psi\rangle = \sum_{i_1 \dots i_n} c_{i_1 \dots i_n} |i_1, \dots, i_n\rangle, \quad (5)$$

with a network of tensors whose dimensions are such that the number of coefficients required to describe the network exhibits a better scaling in n , the number of lattice sites, than does the number of coefficients $c_{i_1 \dots i_n}$ in Eq. (5). Indeed, for many tensor networks this scaling in n is polynomial rather than exponential.

Given this reduction in the number of coefficients in the description, a tensor network ansatz is capable only of representing states which lie within some restricted region of the Hilbert space of the system, but nevertheless these ansätze and associated algorithms are capable of providing substantial insight into the physics of a wide variety of systems in appropriate regimes [e.g., one-dimensional (1D) quantum critical systems, systems with limited long-range entanglement, and 2D systems obeying an area law. Details vary with the specifics of the tensor network employed—see Refs. [20–44] for examples].

In order for a tensor network algorithm to be useful as a tool for numerical computation, it must be possible to perform the operations of the algorithm for a reasonable computational cost. An economical description of the relevant part of the Hilbert space of a system is a good start, but this is not the only factor which must be taken into account: When determining whether a given tensor network algorithm is computationally feasible, the structure of the tensor network itself also plays a significant role. In describing scaling of the cost of a tensor network algorithm, it is customary to express this in the form of a polynomial in some refinement parameter χ , which may (for example) correspond to the dimensions of indices within the tensor network. Assuming that the coefficients of this cost polynomial are small, costs which scale as excessively large powers of χ may then reflect an algorithm which pushes the limits of computational feasibility.

The trade-off between sophistication of an ansatz and the associated cost of the update algorithm represents a critical tension in the development of novel tensor network algorithms. For example, the structure of the 4:1 2D MERA [25,26] indicates that this particular ansatz will provide a powerful representation of highly entangled 2D systems, and this has been confirmed analytically in Ref. [65], where it is shown to furnish a compact and physically meaningful description of the toric code. However, numerical computations using this ansatz are hindered by an update cost of $O(\chi^{26})$. The ability to

FIG. 1. Example graphical representations of tensor networks: (a) matrix multiplication, $C_{ij} = A_{ik} B_{kj}$, and (b) inner product of a matrix and two vectors, $c = v^T M w$.

FIG. 2. Graphical representation of one of the tensor networks which must be contracted during variational optimization of the 3:1 MERA [26,47]. The individual tensors are labeled for subsequent reference, and the tensor expression corresponding to this diagram is given in Eq. (38) of Sec. III.

quickly and conveniently determine the cost of contracting different tensor networks is therefore of great importance to researchers employed in the development of novel tensor network algorithms.

Even when the cost of updating a tensor network algorithm is known, the implementation of these algorithms is frequently a nontrivial affair. The process of contracting a tensor network may always be optimally realized as a series of pairwise contractions (see Appendix B), and the overall cost of contracting the network is highly dependent on the sequence in which these contractions are carried out. For instance, the network shown in Fig. 2 is one network which must be contracted during the variational optimization of the 3:1 1D MERA [26,47]. If each index has dimension χ , then the most efficient contraction sequences yield a cost of $O(\chi^8)$, for example,

$$((((hU)U^*)W^{(1)})W^{(1)*})(W^{(2)}W^{(2)*})), \quad (6)$$

but careless choices of sequence can yield costs as high as $O(\chi^{11})$, for example,

$$(((W^{(2)}U)h)(((W^{(1)}W^{(1)*})U^*)W^{(2)*})). \quad (7)$$

Thus the ability to determine the optimal contraction sequence for a tensor network is as important to those implementing pre-existing tensor network algorithms as it is to those developing them if the algorithms are to be implemented in a computationally efficient manner.

Until recently, manual optimization has largely been the preferred approach within the tensor network community. However, the increasing complexity of tensor network ansätze (e.g., 2D MERA [25–27] and branching MERA [48]) renders this approach increasingly time-consuming and, with an exhaustive search out of the question, it may be difficult to be certain that an optimal contraction sequence has indeed been identified. This situation is exemplified by the 4:1 2D MERA of Refs. [25,26], which was believed for a number of years to have a minimal contraction cost of $O(\chi^{28})$ operations and has only recently been shown to be contractible for only $O(\chi^{26})$ [66]. As tensor network algorithms increase in size and complexity, and implementations begin the transition from single-node

to parallel computing architectures, the limitations of manual optimization are likely to become increasingly apparent.

C. Floating point operations and tensor contractions

When working with tensor network algorithms, the primary objects of interest are single tensor networks made up of a large number of constituent tensors. The principal operation applied to these networks is the contraction of a pair of tensors to form a single tensor, e.g.,

$$C_{ijk} = A_{ik}^l B_{lj}. \quad (8)$$

Evaluation of C_{ijk} proceeds by iteration over indices i , j , and k and summation over index l . If the array representing tensor C_{ijk} is initialized to zero, then construction of C_{ijk} involves $\xi_i \xi_j \xi_k \xi_l$ floating-point multiplications and the same number of additions. With the number of multiplications and additions being equal, it is customary to count only multiplications, and the above calculation is described as having a cost of $\xi_i \xi_j \xi_k \xi_l$ operations.

This may be contrasted with the operation counting adopted by the quantum chemistry and computer science communities. In quantum chemistry ansätze, both tensor contraction and tensor summation play important roles and thus independent tallies must be kept of multiplication and addition operations. The calculation of C_{ijk} given above may in principle be decomposed into two steps [14],

$$(C^*)_{ijkl}^l = A_{ik}^l \times B_{lj} \quad (\text{no sum over } l), \quad (9)$$

$$C_{ijk} = \sum_l (C^*)_{ijkl}^l, \quad (10)$$

where the first calculation is iterated over i , j , k , and l for $\xi_i \xi_j \xi_k \xi_l$ floating-point multiplication operations, and the second is iterated over i , j , and k , and for each set of values $\{i, j, k\}$ the right-hand side is summed over l for a total of $\xi_i \xi_j \xi_k \xi_l$ floating-point addition operations. In practice, there is seldom reason to generate the intermediate object C^* explicitly.

It is noted in Ref. [14] that, following multiplication, it is never suboptimal to immediately sum over all repeated indices on the resulting object (for example, l on C^* in the above example). If this is performed, then the process of multiplication followed by summation over all repeated indices is directly equivalent to contraction of two tensors over all shared indices.

In the present context of optimizing the contraction of a single tensor network, all operations may therefore be understood as pairwise tensor contractions. It is consequently unnecessary to separately count floating-point additions and multiplications (as these numbers are always equal), so for simplicity we shall adopt the approach of the tensor network algorithm community by counting operations which consist of one floating-point multiplication and one floating-point addition apiece. Revisiting the simple examples of Fig. 1, evaluation of C in Eq. (3) therefore incurs a cost of $\xi_i \xi_j \xi_k$ operations, while evaluation of c in Eq. (4) attracts a cost dependent on the order of pairwise contractions. If one begins by contracting v with M ,

$$x^j = v_i M^{ij}, \quad c = x^j w_j, \quad (11)$$

then the cost of computing c is $\xi_i \xi_j + \xi_j$ operations, whereas if one first contracts w with M ,

$$y = M^{ij} w_j, \quad c = v_i y^i, \quad (12)$$

the cost is $\xi_i \xi_j + \xi_i$ operations. Finally, if one begins by performing an outer product between v and w ,

$$z_{ij} = v_i w_j, \quad c = z_{ij} M^{ij}, \quad (13)$$

then this approach incurs a cost of $2\xi_i \xi_j$.

II. SINGLE TERM OPTIMIZATION

This section discusses algorithms for the identification of an optimal (operation-minimizing) contraction sequence for a single-term tensor expression. In describing these algorithms it is assumed that the expression contains no indices of dimension one and that it may not be factorized into two disjoint parts sharing no common indices.

Regarding the first of these requirements, note that an index of dimension one may freely be deleted or reinserted at any time without changing the calculation described, and thus this condition may be imposed without loss of generality.

Regarding the second, note that for any network $\mathcal{N}$ which is factorizable into two or more disjoint subnetworks $\mathcal{N}_1, \mathcal{N}_2, \dots$, e.g.,

$$A^{ij} B^l C_{jk} D_{lm} = (A^{ij} C_{jk})(B^l D_{lm}), \quad (14)$$

and where contraction of these subnetworks yields tensors of dimension greater than 1, an optimal contraction sequence may always be constructed by sequentially concatenating optimal contraction sequences for each subnetwork. In these circumstances we lose no capability (and make substantial performance gains) by addressing each subnetwork individually.¹ Note also that the disjoint or nondisjoint nature of the network is to be assessed only after any indices of dimension one have been deleted.

Finally, it is assumed that no index appears more than once on any of the initial tensors. As mentioned in Ref. [14], it is always optimal to perform the sums over such indices immediately, and thus it is assumed that these have already been performed before the algorithms described in this section are invoked.

¹Where one or more subnetworks does contract to a tensor of dimension one (i.e., a single number, a scalar), the situation is a little more complex. In theory, the optimal time to multiply by this number may arise part-way through the evaluation of another tensor network if the contraction of that network yields an especially small intermediate object. In practice, the best time to multiply by this number is usually during the calculation of human-readable results, which are seldom expressed in terms of large, multi-dimensional tensors. The problem of determining the optimal time to multiply by a scalar is therefore frequently outside the context addressed by the present paper but frequently admits the same approach as discussed above, namely independent contraction of each individual subnetwork.

A. Existing approaches

1. Depth-first constructive approach

We begin with a review of the approach to single-term optimization first described in Ref. [14]. Slightly paraphrased, the algorithm proceeds as follows:

- (1) Compute a single contraction sequence as follows:
 - (a) (i) Let $T_1, \dots, T_n$ be a list of n tensors.
 - (ii) Let $d = n$ be a counter enumerating the total number of unique objects.
 - (iii) Let c be a counter initialized to zero.
 - (b) (i) Increment d .
 - (ii) Choose a pair of tensors $\{T_a, T_b\}$ such that the set indices appearing on T_a is identical to that appearing on T_b , or, if no such pair exists, choose a pair such that $a < b$ and $b > c$.
 - (iii) Contract these two tensors over all common indices (if any). Call the resulting object T_d .
 - (iv) Remove tensors T_a and T_b from the list of tensors.
 - (v) Append T_d to the list of tensors.
 - (vi) Set c to a .
 - (c) Repeat step (1b) until no further actions can be applied. Determine the cost of the contraction sequence performed.
- (2) Repeat step (1) while iterating through all possible choices of pairs $\{T_a, T_b\}$ using a depth-first approach to explore the space of all possible contraction sequences without duplication. Note the cheapest contraction sequence thus found.

Note, in particular, that tensors T_a and T_b need not share any common indices at all. A contraction where tensors do not share any common indices is termed an *outer product*, and there exist tensor networks for which an outer product is a necessary part of the optimal contraction sequence. A simple example is given by

$$D_k = A^i B^j C_{ijk} \quad (15)$$

for $\xi_k > \xi_j$ and $\xi_k > \xi_i$. Writing (XY) to denote the contraction of a tensor X with a tensor Y , the sequence $((AB)C)$ is seen to be cheaper than either $((AC)B)$ or $((BC)A)$.

2. Breadth-first constructive approach

A slightly more sophisticated approach may be found in Ref. [10], and has been incorporated into the TENSOR CONTRACTION ENGINE [12], a tool for optimizing the evaluation of tensor expressions which enjoys widespread use in the quantum chemistry community. Whereas implementation of the depth-first search of Sec. II A 1 will typically be achieved through the use of recursion, the breadth-first approach described in Ref. [10] is iterative and may be summarized as follows:

- (1) Let $S_1 = \{T_1, \dots, T_n\}$ be the set of n tensors which make up network $\mathcal{N}$.
- (2) Let c be a counter running from 2 to n . For each value of c :
 - (a) Let S_c be the set of all objects made up by contracting together c unique tensors from S_1 .

- (b) For each pair of sets S_d, S_{c-d} , $1 \leq d \leq \lfloor \frac{c}{2} \rfloor$, and for each $T_a \in S_d, T_b \in S_{c-d}$ such that each element of S_1 appears at most once in $(T_a T_b)$:
- (i) Determine the cost μ of contracting T_a with T_b .
 - (ii) Where T_a and/or T_b do not belong to S_1 , add to μ the previously determined cost of constructing T_a and/or T_b as appropriate.
 - (iii) Let the contraction sequence Q for constructing this object be written $Q = (T_a T_b)$. Where T_a and/or T_b do not belong to S_1 , optimal contraction sequences for T_a and T_b will have been previously recorded. In Q , replace each appearance of T_a and/or T_b with these optimal contraction sequences.
 - (iv) Locate the object in S_c which corresponds to $(T_a T_b)$. If μ is the cheapest known cost for constructing this object, record the cost μ and the associated contraction sequence Q against this object.
- (3) The optimal cost μ_{best} and a sequence Q_{best} which realizes this are recorded against the only element in S_n .

3. Dynamic programming

Finally, as with the matrix-chain problem discussed in Appendix A, an exhaustive search of all possible contraction sequences may also be performed using dynamic programming with memoization:

- (1) Is network $\mathcal{N}$ trivial (one tensor)? If so, let T denote this tensor, and return zero cost and a contraction sequence $Q = T$.
- (2) Has network $\mathcal{N}$ been costed before? If so, return the known optimal cost and sequence, previously recorded as $\mu_{\text{best}}(\mathcal{N})$ and $Q_{\text{best}}(\mathcal{N})$.
- (3) Otherwise, for each bipartition of $\mathcal{N}$ into two subnetworks $\mathcal{N}_1$ and $\mathcal{N}_2$ (either or both of which may be disjoint):
 - (a) Invoke this algorithm twice more to obtain the optimal cost and sequence for each of $\mathcal{N}_1$ and $\mathcal{N}_2$.
 - (b) Let T_1 be the tensor obtained on fully contracting $\mathcal{N}_1$. Let T_2 be the tensor obtained on fully contracting $\mathcal{N}_2$. Let μ be the cost of the contraction $(T_1 T_2)$.
 - (c) Add to μ the optimal cost for contracting $\mathcal{N}_1$ and the optimal cost for contracting $\mathcal{N}_2$.
 - (d) Define a sequence $Q = (Q_1 Q_2)$ where Q_a is the optimal sequence for contracting subnetwork $\mathcal{N}_a$.
 - (e) If μ is the best cost observed, record the values of μ and Q as $\mu_{\text{best}}(\mathcal{N})$ and $Q_{\text{best}}(\mathcal{N})$, overwriting any previous values thus recorded.
- (4) Return $\mu_{\text{best}}(\mathcal{N})$ and $Q_{\text{best}}(\mathcal{N})$.

Note the role of memoization in preventing unnecessary duplication of efforts: While iterating over all possible bifurcations of networks at all levels of recursion, the same subnetworks will frequently be encountered on many different occasions. As a simple example, given a network $\mathcal{N} = \{ABCDEF\}$, the subnetwork $\{BC\}$ may be encountered by splitting $\mathcal{N}$ into $\{ADEF\}$ and $\{BC\}$ or by splitting $\mathcal{N}$ into $\{ABC\}$ and $\{DEF\}$

and then splitting $\{ABC\}$ into A and $\{BC\}$ or in numerous other ways.

B. Beyond existing approaches

While the approaches described in Sec. II A are adequate for small tensor networks, they become rapidly more expensive as n , the number of tensors in the network, increases. This section introduces two modifications to the search process. The first, described in Sec. II B 1, represents a reordering of the search process in favor of cheapest-first. As the search may now terminate once a contraction sequence has been identified, this approach prunes the search tree to exclude all contraction sequences costing more than the optimal contraction cost of the tensor network. The second modification, described in Sec. II B 2, excludes large numbers of outer product contractions which are shown to be unnecessary for the identification of an optimal contraction sequence.

1. Cost capping

The refinement described in this section may be applied to any of the search algorithms described in Sec. II A. However, to obtain the most benefit it should be applied to an algorithm incorporating some form of memoization, either explicitly (the dynamic programming algorithm of Sec. II A 3) or implicitly (the recording of objects and their associated costs and sequences in the breadth-first algorithm of Sec. II A 2). Of the two, the breadth-first algorithm of Sec. II A 2 is found to perform significantly better with the cost-capping refinement as the initial space to be explored [$n(n-1)/2$ possible pairwise tensor contractions] is both smaller than, and typically contains less expensive (and therefore more relevant) elements than, the initial space of the dynamic programming algorithm (2^{n-1} bifurcations of the initial network $\mathcal{N}$).²

To implement cost capping in the breadth-first algorithm of Sec. II A 2, two modifications are required. First, the sets S_c are initialized empty and are only populated as contraction sequences which yield their elements are identified. Second, pairwise contractions are rejected if their cost exceeds some maximal value μ_{cap} . By choosing a value of μ_{cap} which is sufficiently low, the first invocation of the modified breadth-first search can be guaranteed to terminate with S_n empty, i.e., without constructing an object containing all of the tensors in S_1 . If the value of μ_{cap} is then increased, and the former value is stored in μ_{old} , the existing contents of the sets S_i act as memos for the second invocation of the breadth-first search. On this second invocation it is only necessary to consider contractions where the cost μ of $(T_a T_b)$ satisfies $\mu_{\text{old}} < \mu \leq \mu_{\text{cap}}$ or where $\mu \leq \mu_{\text{old}}$ and one or both of T_a and T_b was itself only constructed during the second invocation. If, following the second invocation, S_n is still empty, μ_{cap} is increased once more, μ_{old} is updated, and the process is repeated. The net result is to yield a cheapest-first construction process where

²Example results for a dynamic programming algorithm with cost capping and some restrictions on outer products may be found in Table V of Ref. [66]. They are most appropriately compared with column “ μ_{cap} ” of Table II in Sec. III, with performance of the breadth-first code being between 8 and 60 times faster.

no object is constructed which costs significantly more than the contraction of network $\mathcal{N}$.

For tensor network algorithms, it is common to define bond dimensions in terms of some refinement parameter χ , assumed to be large. In this context, μ_{cap} may be taken as a bound on the maximum power of χ which appears in the algebraic expression for the contraction cost. Alternatively, one may supply numeric values for the index dimensions and μ_{cap} then comprises a cap on the actual contraction cost.

Let $\xi_{\min}$ denote the dimension of the smallest index in the network, let $\mu_{\text{cap}}^{(i)}$ denote the value of μ_{cap} on invocation i of the algorithm, and let $\mu_{\text{reject}}^{(i)}$ represent the cost of the cheapest rejected pairwise contraction on iteration $i - 1$. We obtained good performance by initializing $\mu_{\text{cap}}^{(1)}$ to 1 and requiring that each subsequent $\mu_{\text{cap}}^{(i)} = \xi_{\min} \mu_{\text{cap}}^{(i-1)}$, except when $\mu_{\text{reject}}^{(i)} > \xi_{\min} \mu_{\text{cap}}^{(i-1)}$, for which we assign $\mu_{\text{cap}}^{(i)} = \mu_{\text{reject}}^{(i-1)}$. However, for networks with many indices of some dimension ξ and a much smaller number of indices of dimension $\xi_{\min} \ll \xi$ we acknowledge that it may be preferable to instigate a more rapid increase in $\xi_{\min}$.

Writing $\text{cost}[(AB)]$ for the cost of contracting tensor A with tensor B , a cost-capped breadth-first algorithm may be realized as follows:

- (1) Let $S_1 = \{T_1, \dots, T_n\}$ be the set of n tensors which make up network $\mathcal{N}$.
- (2) Flag each tensor in S_1 as “old.”
- (3) Let $\{S_i | i \in \mathbb{Z}, 2 \leq i \leq n\}$ be empty sets.
- (4) Let $\mu_{\text{cap}} = 1$, let $\mu_{\text{old}} = 0$, and let $\xi_{\min}$ be the dimension of the smallest index.
- (5) While S_n is empty:
 - (a) Let $\mu_{\text{next}} = \infty$.
 - (b) Let c be a counter running from 2 to n .
For each value of c , and each pair of sets S_d, S_{c-d} , $1 \leq d \leq \lfloor \frac{c}{2} \rfloor$, and each $T_a \in S_d, T_b \in S_{c-d}$ such that each element of S_1 appears at most once in $(T_a T_b)$:
 - (i) Let $\mu = \text{cost}[(T_a T_b)]$.
 - (ii) Where T_a and/or T_b do not belong to S_1 , add to μ the previously determined cost of constructing T_a and/or T_b as appropriate.
 - (iii) If either T_a or T_b is flagged as “new,” let $\mu_0 = 0$. Otherwise, let $\mu_0 = \mu_{\text{old}}$.
 - (iv) If $\mu > \mu_{\text{cap}}$ and $\mu < \mu_{\text{next}}$, let $\mu_{\text{next}} = \mu$.
 - (v) If $\mu_0 < \mu \leq \mu_{\text{cap}}$:
 - (A) Let the contraction sequence $\mathcal{Q}$ for constructing this object be written as $\mathcal{Q} = (T_a T_b)$. Where T_a and/or T_b do not belong to S_1 , the best-known contraction sequences for T_a and T_b will have been previously recorded. In $\mathcal{Q}$, replace each appearance of T_a and/or T_b with these optimal contraction sequences.
 - (B) If no object corresponding to $(T_a T_b)$ has yet been created in S_c , create it. Otherwise, locate the object in S_c which corresponds to $(T_a T_b)$.
 - (C) If μ is the cheapest known cost for constructing this object, then record the

cost μ and the associated contraction sequence $\mathcal{Q}$ against this object and flag the object as “new.”

- (c) Let $\mu_{\text{old}} = \mu_{\text{cap}}$.
- (d) Set μ_{cap} equal to the larger of μ_{next} and $\xi_{\min} \mu_{\text{cap}}$.
- (e) Flag all tensors in all S_i as “old.”
- (6) The optimal cost μ_{best} and a sequence $\mathcal{Q}_{\text{best}}$ which realizes this are recorded against the only element in S_n .

2. Restricting outer products

a. Overview. As noted in Sec. II A 1, there exist tensor networks for which outer products form a necessary part of the optimal contraction sequence. Consequently, when iterating over pairs of tensors to contract in the breadth-first search algorithm, this iteration must include not only pairs of tensors which share one or more common indices but also pairs of tensors which do not. In Secs. II B 2 b–II B 2 c a number of criteria are introduced for the rejection of outer product operations, and it is proven that application of these criteria will never prevent the identification of an optimal contraction sequence. Use is made of the following two lemmas:

Lemma 1 (Combination of shared indices). Suppose there exist two tensors, P and Q , which share two or more indices. Let us denote these indices $\{a_1, a_2, \dots, a_n\}$. We may combine these indices into a single index a .

Proof. Let a enumerate all possible sets of values $\{a_1, a_2, \dots, a_n\}$. ■

Lemma 2 (Combination of external indices). Next, suppose we are only interested in the subnet comprising tensors P and Q . Let us denote all indices on P which do not connect to Q by $\{b_1, b_2, \dots, b_n\}$. We may similarly replace all of these indices by a single combined index b .

Proof. Let b enumerate all possible sets of values $\{b_1, b_2, \dots, b_n\}$. ■

b. Constraints on tensors participating in the outer product. Let A and B be two tensors which we wish to contract together as part of our contraction sequence and which share no common indices. For a nondisjoint network $\mathcal{N}$, object (AB) will necessarily be subsequently contracted with some other tensor C . If C is not a fundamental tensor (i.e., one belonging to S_1), we may always implement our contraction sequence such that C is assembled before contraction (AB) is performed. We will denote the subnetwork comprising tensors A , B , and C by $\mathcal{N}_{ABC}$.

By combining indices in the manner described in Lemmas 1 and 2, the most general form of subnetwork $\mathcal{N}_{ABC}$ is that given in Fig. 3. If an outer product between A and B is a necessary part of the contraction sequence, it follows that sequence $((AB)C)$ is cheaper than either $((AC)B)$ or $((BC)A)$. Comparison with sequence $((AC)B)$ gives us the inequality

$$\begin{aligned} \text{cost}[(AB)C] &< \text{cost}[(AC)B] \\ \xi_a \xi_b \xi_d \xi_e + \xi_a \xi_b \xi_c \xi_d \xi_e &< \xi_a \xi_c \xi_d \xi_e + \xi_a \xi_b \xi_c \xi_e, \end{aligned} \quad (16)$$

which simplifies to

$$\xi_b \xi_d < \xi_c (\xi_b + \xi_d - \xi_b \xi_d). \quad (17)$$

FIG. 3. Two tensors A and B share no common index; should they be combined using an outer product before contracting with tensor C ? Letters $a, \dots, e$ denote indices of dimensions $\xi_a, \dots, \xi_e$, respectively. In Sec. II B 2 b it is shown that when searching for an optimal sequence we need only consider contraction (AB) if indices $a, \dots, e$ satisfy the conditions given in Fig. 4.

Since all dimensions are positive integers, this requires

$$\xi_b + \xi_d - \xi_b \xi_d > 0 \quad (18)$$

and may only be satisfied if $\xi_b = 1$ or $\xi_d = 1$. The condition obtained from sequence $((BC)A)$ is equivalent under the exchanges

$$A \Leftrightarrow B, a \Leftrightarrow b, d \Leftrightarrow e, \quad (19)$$

and so yields

$$\xi_a \xi_e < \xi_c (\xi_a + \xi_e - \xi_a \xi_e), \quad (20)$$

which requires either $\xi_a = 1$ or $\xi_e = 1$. There are, consequently, only three scenarios in which sequence $((AB)C)$ is superior to either $((AC)B)$ or $((BC)A)$, and thus an outer product may be required:

- (1) $\xi_a = \xi_b = 1$: Following an outer product between A and B , all indices on the resulting tensor (AB) are shared with a single tensor C [Fig. 4(a)].
- (2) $\xi_d = \xi_e = 1$: Following an outer product between A and B , the next operation in the contraction sequence is a further outer product between (AB) and C [Fig. 4(b)].
- (3) $\xi_a = \xi_d = 1$ or $\xi_b = \xi_e = 1$: Either A or B is a scalar [Figs. 4(c) and 4(d)]. Provided network $\mathcal{N}$ is nondisjoint, these scenarios will never occur and so may be disregarded.

Regarding scenario 1, substitution of $\xi_a = \xi_b = 1$ into Eqs. (17) and (20) reveals the further constraints

$$\xi_c > \xi_d \quad \xi_c > \xi_e. \quad (21)$$

As scenario (1) is the only admissible scenario in which the outer product object (AB) shares common indices with tensor C , we may conclude that contraction of an outer product

FIG. 4. Scenarios for which the outer product $((AB)C)$ may be of lower cost than either $((AC)B)$ or $((BC)A)$: (a) $\xi_a = \xi_b = 1$; (b) $\xi_d = \xi_e = 1$; (c) $\xi_a = \xi_d = 1$; (d) $\xi_b = \xi_e = 1$.

object (AB) with a tensor C is only necessary if it satisfies the following constraints:

- (1) All indices on tensor (AB) must be shared with tensor C .

- (2) The indices on (AB) and C must satisfy Eq. (21).

Note in particular that when $\xi_c = \xi_d$ or $\xi_c = \xi_e$, the cost of sequence $((AB)C)$ is equal to that of either $((AC)B)$ or $((BC)A)$, respectively, but the latter two sequences do not involve outer products. By discarding sequence $((AB)C)$ when these are equal while retaining either or both of $((AC)B)$ and $((BC)A)$ we guarantee contraction of $\mathcal{N}_{ABC}$ at a cost equal to or less than $\text{cost}[(AB)C]$, without requiring an outer product.

Regarding scenario (2), note that in this context all three sequences $((AB)C)$, $((AC)B)$, and $((BC)A)$ involve outer products. Substitution of $\xi_d = \xi_e = 1$ into Eqs. (17) and (20) yields

$$\xi_c > \xi_a \quad \xi_c > \xi_b, \quad (22)$$

as the conditions under which sequence $((AB)C)$ is superior to sequence $((AC)B)$ and $((BC)A)$. However, if $\text{cost}[(AB)C] = \text{cost}[(AC)B]$, then enforcement of Eq. (22) as a strict inequality will eliminate sequence $((AB)C)$, but (AC) in sequence $((AC)B)$ is also an outer product and under permutation of labels

$$B \Leftrightarrow C, b \Leftrightarrow c \quad (23)$$

this constraint becomes

$$\xi_b > \xi_a \quad \xi_b > \xi_c \quad (24)$$

and sequence $((AC)B)$ will also be rejected. This may be contrasted with scenario (1) where sequence $((AC)B)$ is not an outer product, and is therefore not at risk of being rejected, and hence at least one sequence for contracting $\mathcal{N}_{ABC}$ at a cost equal or less than that of $((AB)C)$ is retained. In order to ensure that the optimal contraction sequence for subnetwork $\mathcal{N}_{ABC}$ is not inadvertently discarded when $\text{cost}[(AB)C]$ is optimal and is equal to $\text{cost}[(AC)B]$ or $\text{cost}[(BC)A]$, the correct constraint to apply when contraction $((AB)C)$ is an outer product is

$$\xi_c \geq \xi_a \quad \xi_c \geq \xi_b, \quad (25)$$

corresponding to the requirement that multiple outer products be performed in nondescending order of tensor dimension.

c. Constraints on tensors contracting with an outer product.

We now consider in further detail the situation where contraction of (AB) with C is not an outer product and tensor C is composite. In Sec. II B 2 b we assumed that a composite tensor C was always constructed before performing the outer product (AB) . If the final two tensors involved in the construction of C are denoted D and E , this corresponds to evaluating (DE) before (AB) in the sequence $((AB)(DE))$. However, with no impact on cost we might equally well choose to evaluate (AB) before (DE) .

Consider now a situation involving composite C where optimal contraction of $\mathcal{N}_{ABC}$ is only achieved through sequence $((AB)C)$ and where the optimal sequence for the construction of C is unique. If contraction of (AB) with (DE) is not an outer product, then by virtue of scenario (1) the most general

FIG. 5. (a) Let tensor C in Fig. 4(a) be composite, with the final step in its construction being the contraction together of two tensors D and E . If contraction (AB) is performed prior to the final step in the construction of C , the most general form of network $\mathcal{N}_{(AB)DE}$ is as shown. (b) If contraction sequence $((AB)(DE))$ is superior to both $((AB)D)E$ and $((AB)E)D$, diagram (a) must reduce to the form shown here (up to an exchange of labels $D \leftrightarrow E$, $a \rightarrow b$, $d \rightarrow e$). (c) Consequently, in scenario (1) of Sec. II B 2 b either tensor C is not composite or the final two tensors in the optimal construction of C admit the form shown here. Note in particular that for a nondisjoint network, tensor C in scenario (1) is consequently never the result of an outer product.

form of a subnetwork comprising tensors (AB) , D , and E is given by Fig. 5(a).

In seeking the optimal contraction sequence for the entire tensor network of which Fig. 5(a) is a part, we only need to consider $((AB)C)$ as a possible optimal sequence for subnetwork $\mathcal{N}_{ABC}$, and thus $((AB)(DE))$ as a possible optimal sequence for $\mathcal{N}_{ABDE}$, if sequence $((AB)(DE))$ is cheaper than both $((AB)D)E$ and $((AB)E)D$. Comparison of $((AB)(DE))$ with $((AB)D)E$ yields the inequality

$$\text{cost}[(AB)(DE)] < \text{cost}[(AB)D)E] \\ \xi_a \xi_b \xi_c \xi_d \xi_e + \xi_a \xi_b \xi_d \xi_e < \xi_a \xi_b \xi_c \xi_d + \xi_b \xi_c \xi_d \xi_e \quad (26)$$

$$\Rightarrow \xi_a \xi_b \xi_d \xi_e < \xi_c (\xi_a \xi_b \xi_d + \xi_b \xi_d \xi_e - \xi_a \xi_b \xi_d \xi_e), \quad (27)$$

and because all index dimensions are positive integers this requires

$$\xi_a \xi_b \xi_d + \xi_b \xi_d \xi_e > \xi_a \xi_b \xi_d \xi_e \\ \Rightarrow \xi_a + \xi_e > \xi_a \xi_e, \quad (28)$$

which can only be satisfied for $\xi_a = 1$ or $\xi_e = 1$. Similarly, requiring

$$\text{cost}[(AB)(DE)] < \text{cost}[(AB)E)D] \\ \xi_a \xi_b \xi_c \xi_d \xi_e + \xi_a \xi_b \xi_d \xi_e < \xi_a \xi_b \xi_c \xi_e + \xi_a \xi_c \xi_d \xi_e \quad (29)$$

yields $\xi_b = 1$ or $\xi_d = 1$.

Letting $\xi_d = \xi_e = 1$ yields a configuration for which the sequence $((AB)(DE))$ is prohibited by scenario (1) of Sec. II B 2 b, and letting $\xi_a = \xi_b = 1$ is prohibited by our assumption that contraction of (AB) with C is not an outer product. Consequently, sequence $((AB)(DE))$ is only preferable to both $((AB)D)E$ and $((AB)E)D$ if either $\xi_a = \xi_d = 1$ or $\xi_b = \xi_e = 1$, both corresponding (up to an interchange of labels) to the network given in Fig. 5(b). It therefore follows that, in addition to the constraints of Sec. II B 2 b, we need only consider contractions $((AB)(DE))$ for (AB) an outer product under the following circumstances:

- (1) when contraction of (AB) with (DE) is itself an outer product or
- (2) when either tensor D or tensor E contributes no unsummed indices to tensor (DE) .

Further, substituting $\xi_b = \xi_e = 1$ back into Eqs. (26) and (29) yields the inequalities

$$\xi_a \xi_c \xi_d + \xi_a \xi_d < \xi_a \xi_c \xi_d + \xi_c \xi_d \Rightarrow \xi_a < \xi_c, \quad (30)$$

$$\xi_a \xi_c \xi_d + \xi_a \xi_d < \xi_a \xi_c + \xi_a \xi_c \xi_d \Rightarrow \xi_d < \xi_c, \quad (31)$$

respectively, which also imply

$$\xi_c^3 > |D|. \quad (32)$$

Given tensors D and E consistent with Fig. 5(c), the constraints of Eqs. (30) and (31) are both necessary and sufficient to ensure the satisfaction of Eqs. (26) and (29), and thus it is only necessary to consider a composite tensor C in scenario (1) of Sec. II B 2 b if these two inequalities are satisfied.

As a slight subtlety, note that contraction of (AB) with E is an outer product, and so for $\xi_d = \xi_c$ the sequence $((AB)E)D$ is *a priori* excluded by Eq. (21). It might therefore appear necessary to relax condition (31) to

$$\xi_d \leq \xi_c, \quad (33)$$

permitting retention of sequence $((AB)(DE))$ as a non-outer-product alternative when $\xi_d = \xi_c$. In practice this situation never arises, as sequences involving the outer product (AB) are never necessary for the optimal contraction of such a network. This result is demonstrated in Appendix C.

We have now shown that if the optimal sequence for constructing tensor C is unique, Fig. 5(c) and Eqs. (30) and (31) define the situations under which it is necessary to consider a composite tensor C in scenario (1) of Sec. II B 2 b. Now consider situations where there exist multiple contraction sequences of optimal cost for the construction of tensor C . If *any* of these sequences are inconsistent with Fig. 5(c) or Eqs. (30) and (31), then there exists a contraction sequence for subnetwork $\mathcal{N}_{ABDE}$ which is cheaper than $((AB)(DE))$, and we may omit consideration of sequences for the full network which involve the contraction of tensor (AB) with tensor (DE) . It is therefore only necessary to consider a composite tensor for the role of tensor C in scenario (1) in Sec. II B 2 b if *no* construction of optimal cost for tensor C exists which is inconsistent with Fig. 5(c), and the *smallest* value of ξ_c encountered in any of these sequences satisfies Eqs. (30) and (31).

Finally, we confirm that if the optimal sequence for tensor C is inconsistent with Fig. 5(c) or Eqs. (30) and (31) but a suboptimal sequence is consistent with these constraints, we still need not consider sequence $((AB)C)$. This follows immediately from the observation that a suboptimal construction for tensor C may never yield the optimal cost for sequence $((AB)C)$.

d. Preferred implementation of these constraints. Reduction of search runtime is best achieved by pruning as many branches of the search tree as possible, as early as possible. To this end, the preferred implementation of these conditions is not only as a restriction on contractions involving tensors which are outer products but also as a restriction on which outer products are performed in the first place. To this end, a

contraction (AB) which is an outer product is only performed if a tensor C is known such that contraction of (AB) with C is not an outer product and

- (1) Subnetwork $\mathcal{N}_{ABC}$ takes the form of Fig. 4(a), with indices satisfying Eq. (21).
- (2) Either tensor C is a fundamental tensor (i.e., belongs to S_1) or the final two tensors in all optimal-cost constructions of C take the form of tensors D and E in Fig. 5(b) with indices satisfying Eqs. (30) and (32).

Note that multiple sequential outer products are not excluded by these constraints, e.g.,

$$(((A_1 A_2) B) C), \quad (34)$$

where both contractions in $((A_1 A_2) B)$ are outer products. For such sequences, if the existence of tensor C satisfies conditions (1) and (2) for the contraction of $(A_1 A_2)$ with B , then it also necessarily satisfies these conditions for the contraction of A_1 with A_2 , even though this contraction is not necessarily immediately followed by contraction of $(A_1 A_2)$ with tensor C .

On the other hand, notice that Eq. (31) is *not* enforced at this time. This is because it is possible for tensor $((A_1 A_2) B)$ to satisfy this inequality while tensor $(A_1 A_2)$ does not. The contraction $(A_1 A_2)$ should not be excluded, however, as it is a necessary precursor to the outer product of tensor $(A_1 A_2)$ with B , and the subsequent contraction $(((A_1 A_2) B) C)$ is then permitted. Enforcement of Eq. (31) is therefore only performed at the time of actually contracting an outer product object with another tensor under scenario (1) of Sec. II B 2 b.

It should be noted that the pre-emptive enforcement of conditions described here augments, not supplants, the enforcement of these constraints during contraction of an outer product object (AB) with a tensor C . The pre-emptive application of some constraints prevents the construction of some unnecessary outer product objects, but the active application of *all* constraints in Sec. II B 2 during the contraction of an outer product object (AB) with another tensor is also necessary to ensure that the outer products which *are* constructed only participate in contractions which are necessary according to the constraints determined above. Such contractions must therefore comply with scenario (1) or (2) of Sec. II B 2 b, along with constraints (21) or (25), respectively. If scenario 1 applies, then tensor C must also either belong to S_1 or have an optimal contraction sequence whose final tensors take the form of Fig. 5(c) of Sec. II B 2 c, with indices satisfying Eqs. (30) and (31).

A pseudocode implementation of all these constraints is rather lengthy and thus is given in Appendix D.

III. RESULTS

This section compares the performance of the basic breadth-first algorithm of Sec. II A 2 with (i) a breadth-first search supplemented by cost capping (Sec. II B 1), (ii) a breadth-first algorithm supplemented by the restrictions on admissible outer products described in Sec. II B 2, and (iii) the full algorithm of Appendix D, which incorporates both cost capping and outer product restrictions. Each of these algorithms was applied to seven tensor networks of varying size and complexity, chosen to be representative of the tensor network contractions

FIG. 6. Sample 1D tensor networks: (a) Application of a time evolution gate in TEBD, corresponding to Eq. (35). (b) Calculation of the environment of a tensor in a 3:1 1D TTN, corresponding to Eq. (36). (c) “Lifting” a term in the Hamiltonian in a 2:1 1D MERA, corresponding to Eq. (39).

commonly encountered in condensed matter physics, where they arise from the following scenarios:

Figure 6(a): Application of a time evolution gate in TEBD. Six tensors. The corresponding expression in index notation is

$$\lambda^{(1)}_{ab} \Gamma^{(1)bd}_c \lambda^{(2)}_{de} \Gamma^{(2)eg}_f \lambda^{(3)}_{gh} U^{cf}_{ij}. \quad (35)$$

Figure 6(b): Variational update of a tensor in a 3:1 1D TTN. Five tensors. The corresponding expression is

$$W^{(2)f}_{cde} h^{kc}_{ab} W^{(1)*ija}_h W^{(2)*bde}_g \rho^{hg}_{lf}. \quad (36)$$

Figure 7(a): Variational update of a tensor in a 9:1 2D TTN. Nine tensors. The corresponding expression is

$$\begin{aligned} & W^{(2)c}_{\varepsilon\zeta\eta\theta\iota\kappa\jmath\lambda\mu} W^{(3)a}_{\text{opqrstuv}} W^{(4)b}_{\text{iwxyz}\alpha\beta\gamma\delta} h^{\chi\jmath\text{hi}}_{\text{mnkl}} \\ & \times W^{(1)*v\xi\pi\rho\sigma\tau\nu\phi\eta}_f W^{(2)*\varepsilon\zeta\eta\theta\iota\kappa\jmath\lambda\mu}_g W^{(3)*\text{opqrstuv}}_d \\ & \times W^{(4)*\text{iwxyz}\alpha\beta\gamma\delta}_e \rho^{fgde}_{\psi cab}. \end{aligned} \quad (37)$$

Figure 2: “Lifting” the Hamiltonian in a 3:1 1D MERA. Seven tensors. The corresponding expression is

$$W^{(1)a}_{efh} W^{(2)b}_{lop} U^{hl}_{im} h^{fi}_{gj} U^{*jm}_{kn} W^{(1)*egk}_c W^{(2)*nop}_d. \quad (38)$$

Figure 6(c): “Lifting” the Hamiltonian in a 2:1 1D MERA. Eleven tensors. The corresponding tensor expression is

$$\begin{aligned} & W^{(1)u}_{ae} W^{(2)v}_{im} W^{(3)w}_{pq} U^{(1)ei}_{dh} U^{(2)mp}_{lo} h^{dhl}_{c g k} \\ & \times U^{(1)*cg}_{bf} U^{(2)*ko}_{jn} W^{(1)*ab}_r W^{(2)*fj}_s W^{(3)*nq}_t. \end{aligned} \quad (39)$$

Figure 7(b): Variational update of an isometry tensor in a 9:1 2D MERA. Nineteen tensors. The corresponding

FIG. 7. Sample 2D tensor networks: (a) Calculation of the environment of a tensor in a 9:1 2D TTN, corresponding to Eq. (37). (b) Calculation of the environment of an isometry tensor in a 9:1 2D MERA, corresponding to Eq. (40). (c) Calculation of the environment of an isometry tensor in a 4:1 2D MERA, corresponding to Eq. (41). For 2D tensor networks, a modified version of the diagrammatic notation is used in which the tensors of a diagram are arranged in layers. Unit cells on the lattice are marked on each layer as a reference guide, and circles containing a downward-pointing arrow represent indices passing to the layer below (drawn to the left), while those containing an upward-pointing arrow represent indices passing to the layer above (drawn to the right). Where a diagram extends over multiple rows, the layer at the left-hand end of one row is immediately above that at the right-hand end of the row below. Circles containing both upward- and downward-pointing arrows indicate either a location on a tensor giving rise to both upward- and downward-going indices or (if not located on a tensor) an index line passing through a given location in a layer. Paired indices are represented by black arrows on a white background, while unpaired indices are represented by a white arrow on black, and the diagrams are periodic, so paired indices departing upwards from the top layer of a diagram enter the bottom layer from beneath and vice versa.

TABLE I. Optimal costs for contraction of the tensor networks given in Eqs. (35)–(41), and example contraction sequences which realize these costs. Sequences were computed using the pure-MATLAB reference implementation of Appendix D [67].

Tensor network	Cost	Example contraction sequence returning output as X
3:1 1D TTN	$4\chi^6$	$X = (((W^{(2)}W^{(2)*})\rho)h)W^{(1)*}$
TEBD	$10\chi^3 + 16\chi^2$	$X = ((\lambda^{(1)}\Gamma^{(1)})(\lambda^{(2)}(\Gamma^{(2)}\lambda^{(3)})))U$
3:1 1D MERA	$2\chi^8 + 2\chi^7 + 2\chi^6$	$X = ((W^{(1)}(U(h(U^*W^{(1)*}))))(W^{(2)}W^{(2)*}))$
9:1 2D TTN	$4\chi^{12} + 4\chi^{10}$	$X = (((W^{(2)}W^{(2)*})(W^{(3)}W^{(3)*})(W^{(4)}W^{(4)*})\rho)))h)W^{(1)*}$
2:1 1D MERA	$2\chi^9 + 4\chi^8 + 2\chi^6 + 2\chi^5$	$X = ((W^{(1)}W^{(1)*})(((W^{(2)}U^{(2)})(U^{(1)}(hU^{(1)*}))(U^{(2)*}W^{(2)*})))W^{(3)}W^{(3)*}))$
9:1 2D MERA	$3\chi^{16} + 3\chi^{14} + \chi^{13}$ $+ \chi^{12} + 5\chi^{10} + 5\chi^9$	$A = (((W^{(2)}W^{(2)*})(V^{(3)}V^{(3)*}))((W^{(3)}W^{(3)*})(W^{(4)}W^{(4)*})(V^{(4)}V^{(4)*})))\rho))$
4:1 2D MERA	$4\chi^{26} + 2\chi^{25} + 2\chi^{23} + 3\chi^{22}$ $+ 3\chi^{20} + \chi^{16} + \chi^{14} + \chi^{13}$ $+ \chi^{12} + 4\chi^8 + 4\chi^7$	$X = (((((A(UU^*))(V^{(2)}V^{(2)*}))(V^{(1)}V^{(1)*}))h)W^{(1)*}))$ $A = ((W^{(2)}W^{(2)*})((((W^{(4)}W^{(4)*})(U^{(1)}(hU^{(1)*})))U^{(3)*})U^{(3)}W^{(5)*})U^{(2)*}))$ $B = (((A(U^{(2)}W^{(5)})(W^{(6)}W^{(6)*}))(W^{(8)}W^{(8)*})(W^{(9)}W^{(9)*})(U^{(4)}U^{(4)*})))$ $X = (((B(W^{(7)}W^{(7)*}))(W^{(3)}W^{(3)*}))\rho)W^{(1)*})$

tensor expression is

$$\begin{aligned}
& W^{(2)\alpha}_{\delta\epsilon\zeta\eta\theta} W^{(3)\beta}_{\iota\kappa\lambda\mu\nu} W^{(4)\gamma}_{\xi\pi\rho\sigma\tau} V^{(1)\bar{\pi}}_{\omega abc} V^{(2)\bar{\xi}}_{\upsilon\phi\chi\psi} \\
& \times V^{(3)\theta\xi}_{defg} V^{(4)\mu\pi}_{hijk} U^{\psi beh}_{lmno} h^{\upsilon\phi\omega l}_{pqrs} U^{*smno}_{tuvw} V^{(1)*rauc}_{xy} \\
& \times V^{(2)*\upsilon q\chi t}_{z\bar{\alpha}} V^{(3)*dvfg}_{\beta\bar{\gamma}} V^{(4)*wijk}_{\delta\bar{\epsilon}} W^{(1)*\bar{\lambda}\bar{\mu}p z x}_{\zeta} \\
& \times W^{(2)*\delta\bar{\alpha}\zeta\eta\bar{\beta}}_{\bar{\eta}} W^{(3)*\gamma\kappa\lambda\bar{\delta}\bar{\nu}}_{\bar{\theta}} W^{(4)*\bar{\gamma}\bar{\epsilon}\rho\sigma\tau}_{\bar{i}} \rho^{\bar{\zeta}\bar{\eta}\bar{\theta}\bar{l}}_{\bar{k}\bar{\alpha}\bar{\beta}\bar{\gamma}}. \quad (40)
\end{aligned}$$

Figure 7(c): Variational update of an isometry tensor in a 4:1 2D MERA. Twenty-seven tensors. The corresponding tensor expression is

$$\begin{aligned}
& W^{(2)\bar{\sigma}}_{\bar{d}\bar{e}\bar{\xi}\bar{\sigma}} W^{(3)\bar{\tau}}_{\bar{f}\bar{g}\bar{\tau}\bar{j}} W^{(4)\bar{\upsilon}}_{\bar{k}\bar{\pi}\bar{m}\bar{\omega}} W^{(5)\bar{\phi}}_{\bar{\rho}\bar{\nu}\bar{a}\bar{f}} W^{(6)\bar{\chi}}_{\bar{\phi}\bar{l}\bar{g}\bar{n}} \\
& \times W^{(7)\bar{\psi}}_{\bar{o}\bar{b}\bar{q}\bar{r}} W^{(8)\bar{\omega}}_{\bar{c}\bar{h}\bar{s}\bar{t}} W^{(9)\bar{a}}_{\bar{i}\bar{p}\bar{u}\bar{v}} U^{(1)\bar{\nu}\bar{\xi}\bar{\pi}\bar{\rho}}_{\bar{\kappa}\bar{\alpha}\bar{\gamma}\bar{\lambda}} U^{(2)\bar{\sigma}\bar{\tau}\bar{\upsilon}\bar{\phi}}_{\bar{\beta}\bar{\chi}\bar{\mu}\bar{\psi}} \\
& \times U^{(3)\bar{\omega}abc}_{\bar{v}\bar{\xi}\bar{d}\bar{e}} U^{(4)\bar{f}ghi}_{\bar{\pi}\bar{j}\bar{k}\bar{l}} h^{\bar{\kappa}\bar{\alpha}\bar{\beta}\bar{\gamma}\bar{\lambda}\bar{\mu}\bar{\nu}\bar{\xi}\bar{\pi}}_{\bar{\delta}\bar{\epsilon}\bar{\zeta}\bar{\eta}\bar{\theta}\bar{\iota}\bar{\kappa}\bar{\lambda}\bar{\mu}} U^{(1)*\bar{\delta}\bar{\epsilon}\bar{\eta}\bar{\theta}}_{\bar{m}\bar{n}\bar{o}\bar{p}} U^{(2)*\bar{\zeta}\bar{\chi}\bar{\iota}\bar{\psi}}_{\bar{s}\bar{t}\bar{u}\bar{v}} \\
& \times U^{(3)*\bar{\kappa}\bar{\lambda}\bar{d}\bar{e}}_{\bar{y}\bar{z}\bar{\alpha}\bar{\beta}} U^{(4)*\bar{\mu}\bar{j}\bar{k}\bar{l}}_{\bar{\zeta}\bar{\eta}\bar{\theta}\bar{\iota}} W^{(1)*\bar{b}\bar{c}\bar{h}\bar{m}}_{\bar{w}} W^{(2)*\bar{d}\bar{e}\bar{n}\bar{s}}_{\bar{x}} \\
& \times W^{(3)*\bar{f}\bar{g}\bar{t}\bar{j}}_{\bar{q}} W^{(4)*\bar{k}\bar{o}\bar{m}\bar{y}}_{\bar{r}} W^{(5)*\bar{p}\bar{u}\bar{z}\bar{z}}_{\bar{w}} W^{(6)*\bar{v}\bar{l}\bar{\eta}\bar{n}}_{\bar{x}} \\
& \times W^{(7)*\bar{o}\bar{a}\bar{q}\bar{r}}_{\bar{\gamma}} W^{(8)*\bar{\beta}\bar{\theta}\bar{s}\bar{t}}_{\bar{\delta}} W^{(9)*\bar{i}\bar{p}\bar{u}\bar{v}}_{\bar{\epsilon}} \rho^{\bar{w}\bar{x}\bar{q}\bar{r}\bar{w}\bar{x}\bar{\gamma}\bar{\delta}\bar{\epsilon}}_{\bar{\rho}\bar{\sigma}\bar{\tau}\bar{\upsilon}\bar{\phi}\bar{\chi}\bar{\psi}\bar{\omega}\bar{a}}. \quad (41)
\end{aligned}$$

All indices were assigned an algebraic dimension χ , except for the physical indices of the TEBD network which were assigned dimension two. Optimal contraction costs and sequences were computed in the limit of large χ (see Appendix E). Due to the number of indices present in these expressions it is necessary to use both Greek and Roman alphabets twice over. The symbol $\bar{\alpha}$ (for example) consequently represents a different index to α . No other meaning is attached to the bar, and no distinction exists between Greek and Roman indices, with all indices ranging from 1 to χ .

All calculations proceeded to completion with the exception of the basic breadth-first algorithm and the breadth-first algorithm with restriction on outer products, neither of which terminated within a reasonable time for Eq. (41). Consequently, these algorithms were only applied to the simpler tensor networks of Eqs. (35)–(40). Aside from these two exceptions, all algorithms successfully identified optimal-cost contraction sequences for all test networks. The optimal costs for contracting Eqs. (35)–(41) and example sequences

realizing these costs are given in Table I, though it should be noted that for most networks there are multiple sequences which realize the optimal cost, and the specific sequence which is returned by an algorithm depends on the precise order in which the search tree is explored. As the descriptions in Sec. II do not specify the order in which an algorithm iterates over the members of a set S_a , the specific sequence which is returned may be implementation dependent, though the optimal cost is not.

The times taken by each algorithm to identify each optimal cost along with a corresponding contraction sequence are given in Table II and plotted in Fig. 8. These results were computed using a Dell workstation with a 3.6-GHz Intel Xeon processor and 48 Gb of 1333-MHz DDR3 RAM, though near-identical performance was obtained on an early-2011 Macbook Pro with a 2.2-GHz Intel i7 processor and 8 Gb of 1333-MHz DDR3 RAM. With cost capping enabled, memory poses no significant constraint, even for networks as large as the 4:1 2D MERA (27 tensors), and from the lack of dependence on processor speed we infer that performance of the algorithm is memory bandwidth limited.

TABLE II. Time in seconds to find a guaranteed optimal contraction sequence for the tensor networks of Figs. 2, 6, and 7 using a breadth-first search. The first column of results corresponds to the basic algorithm of Sec. II A 2, “ μ_{cap} ” indicates application of the cheapest-first algorithm described in Sec. II B 1, and “OP” indicates restriction of allowed contractions creating or involving outer products according to Sec. II B 2.

Tensor network	Number of tensors	Pruning techniques			
		None	OP	μ_{cap}	OP & μ_{cap}
3:1 1D TTN	5	0.0014	0.0014	0.0013	0.0013
TEBD	6	0.0016	0.0014	0.0015	0.0015
3:1 1D MERA	7	0.0025	0.0021	0.0020	0.0019
9:1 2D TTN	9	0.0152	0.0087	0.0036	0.0036
2:1 1D MERA	11	0.0946	0.0086	0.0136	0.0048
9:1 2D MERA	19	7298	1096	0.423	0.069
4:1 2D MERA	27	^a	^a	5507	36

^a Aborted: Insufficient memory to perform calculation without swapping to disk.

FIG. 8. (Color online) Time to determine a guaranteed optimal contraction sequence plotted against number of tensors for different pruning algorithms. Points labeled $\square$ correspond to the exhaustive breadth-first search algorithm of Sec. II A 2 (column “None” in Table II). Points labeled $\circ$ incorporate the restrictions on considered outer products described in Sec. II B 2 (column “OP” in Table II). Points labeled $+$ correspond to the cheapest-first variant on the breadth-first algorithm, described in Sec. II B 1 (column “ μ_{cap} ” in Table II). Points labeled $\times$ correspond to the cheapest-first variant on the breadth-first algorithm with restrictions on outer products (column “OP & μ_{cap} ” in Table II).

The software environment on the workstation comprised MATLAB release 2012b paired with Gnu C++ 4.7.2-ubuntu1 for compilation of the C++ component, while that on the laptop consisted of MATLAB release 2011a paired with Apple Xcode 5.0.2.

From these results the enhanced pruning algorithms of Secs. II B 1 and II B 2 are each seen to be capable of yielding performance increases of up to several orders of magnitude, dependent upon the size and structure of the tensor network being analyzed.

It is recognized that this performance benefit is dependent on the connectivity of the network being studied and that (for example) all benefits of Sec. II B 2 will disappear in the limit of a fully connected network (where every tensor shares an index with every other tensor). However, tensor networks in condensed matter physics are typically much more sparsely connected than the fully connected network, and the MERA networks used as examples here represent a comparatively highly connected example. Furthermore, the benefit persisted even for networks with as few as five or six tensors. The pruning algorithms of Sec. II B are therefore anticipated to be of substantial practical benefit.

IV. DISCUSSION

A. Faster, better ansatz development and implementation

This paper has presented a detailed description (and a reference implementation) of an improved algorithm for determining an optimal sequence for the contraction of a tensor network (or, in the language of quantum chemistry,

a single term in a tensor expression), along with the associated computational cost. The benefit of such automated algorithms is well established in quantum chemistry, with the ability to reduce the development time for novel ansätze by several years [1]. Their adoption in condensed matter physics has been slower primarily due to the larger tensor networks involved, and the NP-hard scaling of such algorithms with the number of tensors. The algorithm presented in this paper is demonstrated to be faster than previously favored exhaustive search algorithms for networks of five or more tensors and is faster by several orders of magnitude for some of the more complex tensor networks encountered in condensed matter physics, making the automated determination of optimal contraction sequences both viable and attractive.

While it could be argued that the tensor network ansätze of condensed matter physics are only just beginning to show the level of complexity where automated determination of optimal contraction sequences is necessary, the advantages of doing so have already been demonstrated: In Ref. [66], one of the authors applied an earlier implementation of some of the techniques described in Sec. II B to the 4:1 2D MERA. At the time, assuming all indices to have a dimension of χ , the best-known contraction sequences for this ansatz scaled as $O(\chi^{28})$ [27]; the automated search revealed a contraction sequence for a cost of $O(\chi^{26})$.

Automated determination of optimal contraction sequences relieves researchers from this lengthy and tiresome task and may on occasion lead to implementations of algorithms which are more efficient than those discovered by hand. As increasingly sophisticated ansätze are proposed, the need for automated search algorithms to determine optimal tensor contraction sequences in condensed matter physics will only increase.

B. New possibilities in condensed matter and quantum gravity

To date, tensor network algorithms in condensed matter physics have employed fixed networks of tensors, where the coefficients of the tensors may be updated but the structure of the network itself remains unchanged, arguably due to the requirement that optimal contraction sequences for these tensor networks be hard-coded into the implementing software at the time of programming. With the advent of an efficient algorithm to find the optimal contraction sequence for an arbitrary network of tensors at runtime, algorithms which involve a dynamically evolving tensor network become feasible. One could, for example, propose an initial ansatz consisting of a very highly connected tensor network with comparatively low index dimensions and then perform variational optimization while allowing the dimensions of these indices to vary according to the spectra of appropriately chosen Schmidt decompositions. By increasing bond dimensions where the spectra are flat and decreasing them where the spectra decay more rapidly, a tensor network may evolve to represent the entanglement structure of the state under study. (A simple example of this is automated reduction of index dimension in DMRG for weakly correlated states, though in this instance the optimal contraction sequence generally remains unchanged.) Use of a software algorithm to determine optimal contraction sequences (and to prohibit changes to the network whose cost is

too great) guarantees that the resulting tensor network ansätze are contracted as efficiently as possible. The concept becomes especially attractive when one considers results presented in Ref. [68] showing how an $O(N)$ speed-up may be obtained for variational optimization of entirely nonsymmetric tensor networks, making self-adaptive networks a potentially viable proposition even in the absence of any obvious physical (e.g., spatial) symmetry.

The Netcon algorithm may also be of interest in the numerical study of loop quantum gravity. In loop quantum gravity, space is represented as a *spin network*, which may be understood as a dynamically evolving network of $SU(2)$ -symmetric tensors, and calculation of observables with respect to a given spin network necessarily involves the contraction of a large tensor network. Further, in evolving from a state A to a state B , multiple tensor contractions may take place, and topological equivalence may give some freedom as to the order in which these contractions are performed. In both situations there is a need to determine the optimal sequence with which to perform tensor contractions, and therefore there is a role for an algorithm such as Netcon. Further, the proven ability of Netcon to analyze networks involving a couple of dozen tensors with only modest requirements in both memory and computation time make the algorithm extremely well suited to calculations on the sort of scales at which current numerical simulations are likely to take place.

ACKNOWLEDGMENTS

R.N.C.P. thanks the Ontario Ministry of Research and Innovation Early Researcher Awards for financial support. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation. F.V. acknowledges support from the European Research Council (Belgium), the Austrian Science Fund (ViCoM), the SFB FoQuS, and the Swiss Institute for Quality Standards. Note on authorship: The cost-capping techniques and restrictions on outer products described in Sec. II B were originally developed by R.N.C.P. in the context of a depth-first search as described in Sec. II A 1. J.H. and F.V. then implemented cost capping in the context of a breadth-first search as per Sec. II A 2 and demonstrated that this yielded a further substantial improvement in performance. The authors combined their developments and proceeded collaboratively to publication.

APPENDIX A: MATRIX CHAIN MULTIPLICATION

The matrix chain multiplication problem is a well-known problem in linear algebra, in which one must determine the least number of mathematical operations required to evaluate the product of a series of varyingly sized matrices. For example, given the calculation

$$D_{ij} = A_{ik} \times B_{kl} \times C_{lj}, \quad (\text{A1})$$

where indices i, j, k , and l range from 1 to ξ_i, ξ_j, ξ_k , and ξ_l , respectively (and ξ_x is termed the *dimension* of index x), one may first multiply A by B and then multiply (AB) by C at a cost of $\xi_i \xi_k \xi_l + \xi_i \xi_l \xi_j$ operations or first multiply B by C and then multiply A by (BC) at a cost of $\xi_k \xi_l \xi_j + \xi_i \xi_k \xi_j$

operations, where each operation comprises one floating point multiplication followed by one floating point addition.

A matrix may only ever be multiplied by one of at most two immediate neighbours,³ and, as a consequence, an optimal sequence of pairwise matrix products may always be found in polynomial time scaling as $O(n^3)$ using dynamic programming techniques [69]. For networks where a tensor may share indices with more than two immediate neighbors, however, this reduction to polynomial time breaks down, and, as seen in Sec. II B 2, the optimal contraction sequence may then even require performing the outer product between two or more tensors which do not share a common index. (Constraints derived in Sec. II B 2 b show that there must exist at least one tensor sharing indices with three neighbors for this to be necessary, and so outer products are never required for the matrix chain.) One is therefore forced to consider all pairwise contractions regardless of whether any indices are shared by the participating tensors, and the problem of determining an optimal contraction sequence for a general tensor network has been shown to be NP-hard [14].

For matrix chain multiplication, an even faster algorithm is known which returns an optimal sequence in time scaling as $O(n \log n)$ [70,71]. However, no extension of this algorithm to more general tensor networks has yet been proposed.

APPENDIX B: PREFERENTIAL NATURE OF PAIRWISE CONTRACTIONS

In Sec. I B it was stated that an optimal contraction sequence for a tensor network may always be realized as a series of pairwise contractions. To see this, consider the contraction of three tensors, A , B , and C , to yield a single tensor D . Let ξ_{AB} denote the product of the dimensions of all indices on tensor A which connect to tensor B , and similarly for ξ_{AC} and ξ_{BC} . Let ξ_A denote the product of all indices on tensor A which do not connect to either B or C , and similarly for ξ_B and ξ_C . The dimension of a set containing no indices is always one. Contracting these three tensors as a single process involves a cost of

$$2\xi_A \xi_B \xi_C \xi_{AB} \xi_{AC} \xi_{BC}. \quad (\text{B1})$$

For example, for the contraction

$$D_e^a = A^{abc} B_{bd} C_{ce}^d, \quad (\text{B2})$$

we have

$$\begin{aligned} \xi_A &= |a|, \quad \xi_B = 1, \quad \xi_C = |e|, \\ \xi_{AB} &= |b|, \quad \xi_{AC} = |c|, \quad \text{and } \xi_{BC} = |d|. \end{aligned} \quad (\text{B3})$$

For each element of D_e^a it is necessary to sum over $\xi_{AB} \xi_{AC} \xi_{BC}$ different contributions (corresponding to the enumeration of indices b, c , and d), each involving two multiplications, and there are then $\xi_A \xi_B \xi_C$ entries in D_e^a , for the total number of multiplication operations given in Eq. (B1). In contrast, pairwise contraction may be achieved by any of the

³Indeed, the constraints of Sec. II B 2 confirm that outer products need never be considered when looking for the optimal contraction of a matrix chain.

sequences $((AB)C)$, $((AC)B)$, or $((BC)A)$, where $((XY)Z)$ means “contract tensor X with tensor Y , then contract the result with tensor Z .” The sequence $((AB)C)$ is readily seen to attract a total cost of

$$\xi_A \xi_B \xi_{AB} \xi_{AC} \xi_{BC} + \xi_A \xi_B \xi_C \xi_{AC} \xi_{BC} \quad (\text{B4})$$

multiplication operations, with those for $((AC)B)$ and $((BC)A)$ being achieved by the relevant label permutations. Since all parameters in Eqs. (B1) and (B4) take a value greater than or equal to 1, the value of Eq. (B4) is always less than or equal to that of Eq. (B1). The argument extends directly to contraction of four or more tensors, with the cost of sequential pairwise contraction continuing to always be less than or equal to that of more complicated contractions, and, consequently, for any tensor network it is always possible to identify a minimum-cost contraction sequence in which all contractions proceed in a pairwise fashion. Note that no assumption has been made about the values of the index dimensions, and thus this result holds even for contraction sequences involving outer products, which may be equated with contraction over indices of dimension one.

APPENDIX C: EXCLUSION OF SEQUENCES $((AB)(DE))$ AND $((AB)E)D$ FOR $\xi_d = \xi_c$ IN FIG. 5(c)

In Sec. II B 2 c we showed that to find an optimal contraction sequence we only need consider outer products where the resulting object is contracted with a fundamental tensor (supplied as input to the algorithm), or with a composite tensor (denoted C) where the final step in the lowest-cost construction of this tensor necessarily takes the form of contracting D with E in Fig. 5(c). It was also shown that indices ξ_a , ξ_c , and ξ_d must satisfy the constraints

$$\xi_d < \xi_c \quad (30)$$

$$\xi_d \leq \xi_c, \quad (33)$$

and that where multiple lowest-cost constructions of tensor C exist, all must be consistent with Fig. 5(c) and it is the smallest value of ξ_c obtained which must satisfy Eqs. (30) and (33). We now show that Eq. (33) may be tightened to

$$\xi_d < \xi_c \quad (31)$$

In Sec. II B 2 c we considered the contraction of a tensor (AB) , formed by performing an outer product between tensors A and B , with tensors D and E which make up composite tensor C , under the assumption that contraction (AB) was always performed first. We then asked when it was preferable to contract tensor D with tensor E , thus obtaining tensor C , rather than contracting tensor (AB) directly with either D or E . We now relax this setup to consider the network of four tensors, A , B , D , and E , prior to any contraction but subject to the constraint that $\xi_c = \xi_d$. This network is shown in Fig. 9, and Eq. (30) now corresponds to

$$\xi_{a_1} \xi_{a_2} < \xi_c, \quad (\text{C1})$$

where $\xi_{a_1} \xi_{a_2} = \xi_a$.

From Sec. II B 2 c we know that when $\xi_c = \xi_d$ the costs of sequences $((AB)(DE))$ and $((AB)E)D$ are equal and are

FIG. 9. Network $\mathcal{N}_{ABDE}$ of Fig. 5(b) prior to contraction of tensor A with tensor B .

given by

$$\xi_{a_1} \xi_{a_2} + \xi_{a_1} \xi_{a_2} \xi_c^2 + \xi_{a_1} \xi_{a_2} \xi_c. \quad (\text{C2})$$

We now compare this with the cost of sequence $((AE)D)B$, which evaluates to

$$\xi_{a_1} \xi_c + \xi_{a_1} \xi_{a_2} \xi_c^2 + \xi_{a_2} \xi_c. \quad (\text{C3})$$

For $((AB)(DE))$ or $((AB)E)D$ to be cheaper than $((AE)D)B$ requires

$$\begin{aligned} \xi_{a_1} \xi_{a_2} + \xi_{a_1} \xi_{a_2} \xi_c^2 + \xi_{a_1} \xi_{a_2} \xi_c &< \xi_{a_1} \xi_c + \xi_{a_1} \xi_{a_2} \xi_c^2 + \xi_{a_2} \xi_c \\ \xi_{a_1} \xi_{a_2} &< (\xi_{a_1} + \xi_{a_2} - \xi_{a_1} \xi_{a_2}) \xi_c \\ \Rightarrow 0 &< \xi_{a_1} + \xi_{a_2} - \xi_{a_1} \xi_{a_2}. \end{aligned} \quad (\text{C4})$$

Nondisjointness of the network requires $\xi_{a_1} \geq 2$ and $\xi_{a_2} \geq 2$, so this condition is never satisfied, and we may be sure of identifying an optimal contraction sequence without considering either $((AB)(DE))$ or $((AB)E)D$. Consequently, it is entirely acceptable that $((AB)E)D$ be excluded by Eq. (21) and $((AB)(DE))$ by Eq. (31). Meanwhile, for sequence $((AE)D)B$ the constraints of Eq. (21) become

$$\xi_{a_1} \xi_c > \xi_{a_2} \quad \xi_{a_1} \xi_c > \xi_c \quad (\text{C5})$$

and are automatically satisfied by virtue of Eq. (C1) and $\xi_{a_1} \geq 2$.

APPENDIX D: PSEUDOCODE IMPLEMENTATION OF SEC. II B

This Appendix presents a pseudocode implementation of the pruning algorithms described in Sec. II B 1 and Sec. II B 2. The constraints on tensor C of Sec. II B 2 c are realized by constructing a list (called L) containing the sets of index labels appearing on admissible tensors, along with the index dimension corresponding to $\xi_c = |E|$ in Fig. 5(c). Whenever the best-known contraction sequence for a tensor C ends with a contraction (DE) consistent with the structure shown in Fig. 5(c), its details are added to the list. However, if a sequence is subsequently identified of equal or lower cost which does not satisfy this construction, the details of C are removed from the list again. If this is going to happen, it always does so before the next increment of μ_{cap} . Additions to the list are therefore made provisionally, to be confirmed once a pass with a given value of μ_{cap} is complete, and outer products are only performed if there exists at least one nonprovisional entry in the list which is consistent with the requirements for tensor C given in Sec. II B 2 c.

If any provisional entries are present on the list when μ_{cap} is due to be increased, then these entries are confirmed and another pass is performed at the same value of μ_{cap} to resolve the new contractions which are made possible by the additional nonprovisional entries in the list. Waiting until the initial pass at cost μ_{cap} is complete before performing the newly available outer products means that if a provisional entry is added to the list but is then subsequently deleted, or its constraints are tightened, exploration of the corresponding unnecessary branches of the search tree can largely be avoided.⁴

As per Sec. II it is assumed that all indices are of dimension 2 or higher and that the tensor network is nondisjoint. A discussion of how this algorithm may be extended to disjoint networks and networks including indices of dimension 1 may be found in the instructions accompanying the reference implementation [67].

Algorithm: NETCON

- (1) Let $S_1 = \{T_1, \dots, T_n\}$ be the set of n tensors which make up network $\mathcal{N}$.
- (2) Flag each tensor in S_1 as “old.”
- (3) Let $\{S_i | i \in \mathbb{Z}, 2 \leq i \leq n\}$ be empty sets.
- (4) Let $\mu_{\text{cap}} = 1$, let $\mu_{\text{old}} = 0$, let $\mu_{\text{next}} = \infty$, and let $\xi_{\min}$ be the dimension of the smallest index on any tensor.
- (5) Let L be an empty list whose entries $l \in L$ each take the form $l \equiv \{I_l, \xi_l, f_l\}$, where I_l is a list of indices, ξ_l is a tensor dimension, and f_l is a numerical flag. For each T_i in turn, perform subalgorithm ADDTOLIST $[T_i, -]$.
- (6) Assign flag 0 (“old entry”) to all entries in L .
- (7) While S_n is empty:
 - (a) Let c be a counter running from 2 to n .
For each value of c ,
and each pair of sets S_d, S_{c-d} , $1 \leq d \leq \lfloor \frac{c}{2} \rfloor$,
and each $T_a \in S_d, T_b \in S_{c-d}$ such that each element of S_1 appears at most once in $(T_a T_b)$:
 - (i) If T_a and T_b share no common indices:
 - (A) If $\mu_{\text{old}} \neq \mu_{\text{cap}}$, and there exists an entry l in L with flag $f_l = 0$ (“old”) for which I_l contains all indices on T_a and T_b and $\xi_l \geq |T_a||T_b|$, advance to step (7(a)iv).
 - (B) If $\mu_{\text{old}} = \mu_{\text{cap}}$, and there exists an entry l in L with flag $f_l = 1$ (“new”) for which I_l contains all indices on T_a and T_b and $\xi_l \geq |T_a||T_b|$, advance to step (7(a)iv).

⁴This avoidance may not be perfect; consider the following situation: Let a tensor X , consistent with (DE) in Fig. 5(c), be added to list L and, subsequently, become nonprovisional. Let another tensor Y be added to list L and become nonprovisional at the same time or later than X . If addition of tensor Y to list L makes it possible to construct X for equal-best or better cost than previously known, and in a manner either inconsistent with Fig. 5(c) or yielding a lower value of ξ_c than previously obtained, then it is possible that time could have been wasted on the performance of unnecessary contractions (AB) and $((AB)X)$, where (AB) is an outer product, prior to finding the new sequence for X and implementing the improved constraints which that sequence implies. In practice, this scenario has yet to be observed.

- (C) If $\mu_{\text{old}} = \mu_{\text{cap}}$, either of T_a or T_b is flagged as “new,” and there exists an entry l in L with flag $f_l = 0$ for which I_l contains all indices on T_a and T_b and $\xi_l \geq |T_a||T_b|$, advance to step (7(a)iv).
 - (D) Otherwise, return to step (7a) and select the next pair $\{T_a, T_b\}$.
- (ii) If either T_a or T_b is the result of an outer product, and contraction of T_a with T_b is not an outer product:
 - (A) If both T_a and T_b are outer products, this violates the constraints of Sec. IIB2c. Return to step (7a) and select the next pair $\{T_a, T_b\}$.
 - (B) Let Y be the member of $\{T_a, T_b\}$ which is an outer product, and X be the member of $\{T_a, T_b\}$ which is not an outer product. If X is not in S_1 and does not satisfy the form mandated in Fig. 5(c), return to step (7a) and select the next pair $\{T_a, T_b\}$.
 - (C) If X and Y do not satisfy Eq. (21), or X is not in S_1 and X and Y do not satisfy Eqs. (30) and (31), return to step (7a) and select the next pair $\{T_a, T_b\}$.
 - (iii) If either T_a or T_b is the result of an outer product, and contraction of T_a with T_b is an outer product, check that it satisfies Eq. (25). Otherwise, return to step (7a) and select the next pair $\{T_a, T_b\}$.
 - (iv) Let $\mu = \text{cost}[(T_a T_b)]$.
 - (v) Where T_a and/or T_b do not belong to S_1 , add to μ the previously determined cost of constructing T_a and/or T_b as appropriate.
 - (vi) If either T_a or T_b is flagged as “new,” let $\mu_0 = 0$. Otherwise, let $\mu_0 = \mu_{\text{old}}$.
 - (vii) If $\mu > \mu_{\text{cap}}$ and $\mu < \mu_{\text{next}}$, let $\mu_{\text{next}} = \mu$.
 - (viii) If $\mu_0 < \mu \leq \mu_{\text{cap}}$:
 - (A) Let the contraction sequence $\mathcal{Q}$ for constructing this object be written $\mathcal{Q} = (T_a T_b)$. Where T_a and/or T_b do not belong to S_1 , the best-known contraction sequences for T_a and T_b will have been previously recorded: In $\mathcal{Q}$, replace each appearance of T_a and/or T_b with the corresponding best-known contraction sequences.
 - (B) If no object corresponding to $(T_a T_b)$ has yet been created in S_c , create it. Otherwise, locate the object in S_c which corresponds to $(T_a T_b)$.
 - (C) If this is the first known sequence for constructing this object, or μ is cheaper than any previously known cost for constructing this object:
 - (I) Record the cost μ and the associated contraction sequence $\mathcal{Q}$ against this object.
 - (II) Flag the object as “new.”

- (III) If T_a and T_b have structures permitting them to be identified with D and E in Fig. 5(c) and their dimensions satisfy Eq. (32), perform sub-algorithm ADDTOLIST $[(T_a T_b), Q]$. Otherwise, if this is not the first known sequence for constructing this object, perform subalgorithm REMOVEFROMLIST $[(T_a T_b)]$.
- (D) If this is not the first known sequence for constructing this object, μ is equal to the best previously known cost for constructing this object, and the final contraction in the previous best-known sequence for this object can be identified with (DE) in Fig. 5(c):
- (I) If T_a and T_b have structures permitting them to be identified with D and E in Fig. 5(c), and the value of ξ_c in Fig. 5(c) for sequence Q is lower than that of the previous best recorded sequence for $(T_a T_b)$ but still satisfies Eq. (32), replace that sequence with Q and then perform subalgorithm UPDATERLIST $[(T_a T_b), Q]$.
 - (II) If T_a and T_b have structures permitting them to be identified with D and E in Fig. 5(c), and the value of ξ_c in Fig. 5(c) for sequence Q is lower than that of the previous best recorded sequence for $(T_a T_b)$ and does not satisfy Eq. (32), replace that sequence with Q and then perform subalgorithm REMOVEFROMLIST $[(T_a T_b), Q]$.
 - (III) If T_a and T_b do not have structures permitting them to be identified with D and E in Fig. 5(c), replace the previous best recorded sequence for $(T_a T_b)$ with Q and perform subalgorithm REMOVEFROMLIST $[(T_a T_b)]$.
- (E) If $(T_a T_b) \in S_n$, set $\mu_{\text{cap}} = \mu$. (Solution found; no need to consider any sequences more expensive than this.)
- (b) Let $\mu_{\text{old}} = \mu_{\text{cap}}$.
 - (c) If no entries in L are flagged 2 (“provisional”), set μ_{cap} equal to the larger of μ_{next} and $\xi_{\min} \mu_{\text{cap}}$ and then set $\mu_{\text{next}} = \infty$.
 - (d) Flag all tensors in all S_i as “old.”
 - (e) Confirm provisional additions to L : Decrease the flags on all entries in L by 1 to a minimum of 0.
 - (f) Remove redundant entries in L : For each entry $l \in L$ and for each entry $l' \in \{L | f_{l'} = 1, l' \neq l\}$, if $I_l \subset I_{l'}$ and $\xi_l \leq \xi_{l'}$, delete entry l from L .
- (8) The optimal cost μ_{best} and a sequence Q_{best} which realizes this are recorded against the only element in S_n .
- Subalgorithm: ADDTOLIST $[X, Q]$
- (1) If X is in S_1 , let $\xi = \infty$. Otherwise, considering the final step in sequence Q , let ξ be the dimension of the tensor which corresponds to E in Fig. 5(c).
 - (2) An entry in L takes the form $l = \{I_l, \xi_l, f_l\}$ where I_l is a list of indices, ξ_l is a tensor dimension, and f_l is a flag. Stepping through all nonprovisional entries $\{l \in L | f_l \neq 2\}$:
 - (a) If indices on X exactly match I_l :
 - (i) Set $\xi_l = \xi$.
 - (ii) If this corresponds to an increase in the value of ξ_l , constraints have been relaxed. Set $f_l = 2$ to ensure another pass at the same value of μ_{cap} .
 - (iii) Terminate subalgorithm ADDTOLIST.
 - (b) If $f_l \neq 2$, all indices on X appear in I_l , and $\xi \leq \xi_l$:
 - (i) An entry for X would be redundant. One may previously have been created in L due to an earlier sequence with a higher value of ξ . If so, this is now superceded by the present value of ξ and so is also redundant. Perform subalgorithm REMOVEFROMLIST $[X]$.
 - (ii) Terminate subalgorithm ADDTOLIST.
 - (3) If an entry for X may previously have been created in L with higher ξ , step through all provisional entries $\{l \in L | f_l = 2\}$. If an entry l is found for which I_l exactly matches the indices on X :
 - (a) Set $\xi_l = \xi$.
 - (b) Terminate subalgorithm ADDTOLIST.
 - (4) Add a new entry l' to L , where $I_{l'}$ is a list of all indices appearing on X , $\xi_{l'} = \xi$, and $f_{l'} = 2$ (“provisional”).
- Subalgorithm: UPDATERLIST $[X, Q]$
- (1) Considering the final step in sequence Q , let ξ be the dimension of the tensor which corresponds to E in Fig. 5(c).
 - (2) An entry in L takes the form $l = \{I_l, \xi_l, f_l\}$, where I_l is a list of indices, ξ_l is a tensor dimension, and f_l is a flag. Stepping through all entries $\{l \in L\}$, if I_l exactly matches the indices on X :
 - (a) Set $\xi_l = \xi$.
 - (b) Terminate subalgorithm UPDATERLIST.
- Subalgorithm: REMOVEFROMLIST $[X]$
- (1) An entry in L takes the form $l = \{I_l, \xi_l, f_l\}$, where I_l is a list of indices, ξ_l is a tensor dimension, and f_l is a flag. Stepping through all entries $l \in L$, if the indices on X exactly match I_l :
 - (a) Delete entry l from L .
 - (b) Terminate subalgorithm REMOVEFROMLIST.
- A reference implementation of the NETCON algorithm, written in MATLAB and C++, may be found in the Supplemental Material associated with this paper [67]. The implementation requires MATLAB 2011a or above, and performance may be enhanced using a compatible C++ compiler. It has been tested under MATLAB 2011a with Apple XCode 5.0.2 and MATLAB 2012b with Gnu C++ 4.7.2-ubuntu1. The network contraction sequences output by the reference implementation are fully compatible with the tensor network contraction

packages `ncon()` and `multienv()` of Refs. [72] and [68], respectively.

APPENDIX E: OPTIMAL CONTRACTION COST AND INDEX DIMENSION

In Sec. III we computed contraction sequences for seven different tensor networks, expressing the costs of these sequences as polynomials in a parameter χ corresponding to index dimension. The costs and sequences computed are optimal for sufficiently large values of χ , but it is important to note that the optimal contraction sequence may depend on the value of χ in a nontrivial fashion. As an example, consider the simple network shown in Fig. 10. Contracting this network according to the sequence $((AC)B)$ incurs a cost of $72\chi^6$, while the sequence $((BC)A)$ incurs a cost of $12\chi^7$. The sequence $((AC)B)$ is therefore clearly to be preferred in the limit of large χ , but sequence $((BC)A)$ is preferred for index dimensions corresponding to values of $\chi < 6$. When comparing the costs of different contraction sequences it is therefore necessary either to know the value of χ or to assume that χ is large. The

FIG. 10. A simple network for which the optimal contraction sequence varies with the value of χ .

reference implementation of Appendix D which is provided in the Supplemental Material [67] permits either scenario, allowing index dimensions to be specified either as explicit real numbers or as monomials $a\chi^b$ for real positive a and real non-negative integer b . When index dimensions are specified as monomials, the value of χ is assumed to be large.

- [1] A. A. Auer, G. Baumgartner, D. E. Bernholdt, A. Bibireata, V. Choppella, D. Cociorva, X. Gao, R. Harrison, S. Krishnamoorthy, S. Krishnan, C.-C. Lam, Q. Lu, M. Nooijen, R. Pitzer, J. Ramanujam, P. Sadayappan, and A. Sibiryakov, *Mol. Phys.* **104**, 211 (2006).
- [2] G. Baumgartner, D. E. Bernholdt, D. Cociorva, C.-C. Lam, J. Ramanujam, R. Harrison, M. Nooijen, and P. Sadayappan, in *Proceedings of the 16th International Parallel and Distributed Processing Symposium, IPDPS '02* (IEEE Computer Society, Washington, DC, 2002), p. 33.
- [3] G. Baumgartner, D. Bernholdt, D. Cociorva, R. Harrison, S. Hirata, C. Lam, M. Nooijen, R. Pitzer, J. Ramanujam, and P. Sadayappan, in *Supercomputing, ACM/IEEE 2002 Conference* (IEEE, Baltimore, Maryland, 2002), p. 5.
- [4] G. Baumgartner, A. Auer, D. Bernholdt, A. Bibireata, V. Choppella, D. Cociorva, X. Gao, R. Harrison, S. Hirata, S. Krishnamoorthy, S. Krishnan, C. Lam, Q. Lu, M. Nooijen, R. Pitzer, J. Ramanujam, P. Sadayappan, and A. Sibiryakov, *Proc. IEEE* **93**, 276 (2005).
- [5] D. Cociorva, J. Wilkins, G. Baumgartner, P. Sadayappan, J. Ramanujam, M. Nooijen, D. Bernholdt, and R. Harrison, in *High Performance Computing—HiPC 2001*, Lecture Notes in Computer Science Vol. 2228, edited by B. Monien, V. Prasanna, and S. Vajapeyam (Springer, Berlin, 2001), pp. 237–248.
- [6] D. Cociorva, J. W. Wilkins, C. Lam, G. Baumgartner, J. Ramanujam, and P. Sadayappan, in *Proceedings of the 15th International Conference on Supercomputing, ICS '01* (ACM, New York, 2001), pp. 103–113.
- [7] D. Cociorva, G. Baumgartner, C.-C. Lam, P. Sadayappan, J. Ramanujam, M. Nooijen, D. E. Bernholdt, and R. Harrison, in *Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation, PLDI '02* (ACM, New York, 2002), pp. 177–186.
- [8] D. Cociorva, G. Baumgartner, C.-C. Lam, P. Sadayappan, J. Ramanujam, M. Nooijen, D. E. Bernholdt, and R. Harrison, *SIGPLAN Not.* **37**, 177 (2002).
- [9] D. Cociorva, X. Gao, S. Krishnan, G. Baumgartner, C. Lam, P. Sadayappan, and J. Ramanujam, *Parallel and Distributed Processing Symposium, 2003. Proceedings International* (IEEE, Baltimore, Maryland, 2003).
- [10] A. Hartono, A. Sibiryakov, M. Nooijen, G. Baumgartner, D. Bernholdt, S. Hirata, C.-C. Lam, R. Pitzer, J. Ramanujam, and P. Sadayappan, in *Computational Science—ICCS 2005*, Lecture Notes in Computer Science Vol. 3514, edited by V. Sunderam, G. Albada, P. Sloot, and J. Dongarra (Springer, Berlin, 2005), pp. 155–164.
- [11] A. Hartono, Q. Lu, X. Gao, S. Krishnamoorthy, M. Nooijen, G. Baumgartner, D. Bernholdt, V. Choppella, R. Pitzer, J. Ramanujam, A. Rountev, and P. Sadayappan, in *Computational Science—ICCS 2006*, Lecture Notes in Computer Science Vol. 3991, edited by V. Alexandrov, G. Albada, P. Sloot, and J. Dongarra (Springer, Berlin, 2006), pp. 267–275.
- [12] S. Hirata, *J. Phys. Chem. A* **107**, 9887 (2003).
- [13] S. Hirata, P.-D. Fan, A. A. Auer, M. Nooijen, and P. Piecuch, *J. Chem. Phys.* **121**, 12197 (2004).
- [14] C.-C. Lam, P. Sadayappan, and R. Wenger, *Parallel Process. Lett.* **07**, 157 (1997).
- [15] C.-C. Lam, P. Sadayappan, and R. Wenger, in *Proceedings of Eighth SIAM Conference on Parallel Processing for Scientific Computing* (SIAM, New York, 1997).
- [16] C. Lam, P. Sadayappan, D. Cociorva, M. Alouani, and J. Wilkins, in *Ninth SIAM Conference on Parallel Processing for Scientific Computing* (SIAM, New York, 1999).
- [17] C.-C. Lam, D. Cociorva, G. Baumgartner, and P. Sadayappan, in *Languages and Compilers for Parallel Computing*, Lecture

- Notes in Computer Science Vol. 1863, edited by L. Carter and J. Ferrante (Springer, Berlin, 2000), pp. 350–364.
- [18] R. Penrose, in *Quantum Theory and Beyond*, edited by T. Bastin (Cambridge University Press, Cambridge, 1971).
 - [19] R. N. C. Pfeifer, Ph.D. thesis, The University of Queensland, 2011, [arXiv:1202.1522v2](#) [cond-mat.str-el].
 - [20] A. Feiguin, S. Trebst, A. W. W. Ludwig, M. Troyer, A. Kitaev, Z. Wang, and M. H. Freedman, *Phys. Rev. Lett.* **98**, 160409 (2007).
 - [21] S. Trebst, E. Ardonne, A. Feiguin, D. A. Huse, A. W. W. Ludwig, and M. Troyer, *Phys. Rev. Lett.* **101**, 050401 (2008).
 - [22] R. König and E. Bilgin, *Phys. Rev. B* **82**, 125118 (2010).
 - [23] R. N. C. Pfeifer, G. Evenbly, and G. Vidal, *Phys. Rev. A* **79**, 040301(R) (2009).
 - [24] R. N. C. Pfeifer, P. Corboz, O. Buerschaper, M. Aguado, M. Troyer, and G. Vidal, *Phys. Rev. B* **82**, 115126 (2010).
 - [25] L. Cincio, J. Dziarmaga, and M. M. Rams, *Phys. Rev. Lett.* **100**, 240603 (2008).
 - [26] G. Evenbly and G. Vidal, *Phys. Rev. B* **79**, 144108 (2009).
 - [27] G. Evenbly and G. Vidal, *Phys. Rev. Lett.* **102**, 180406 (2009).
 - [28] G. Evenbly, R. N. C. Pfeifer, V. Picó, S. Iblisdir, L. Tagliacozzo, I. P. McCulloch, and G. Vidal, *Phys. Rev. B* **82**, 161107 (2010).
 - [29] R. M. Noack and S. R. White, *Phys. Rev. B* **47**, 9243 (1993).
 - [30] M. Fannes, B. Nachtergaele, and R. F. Werner, *J. Stat. Phys.* **66**, 939 (1992).
 - [31] B. Friedman, *J. Phys.: Condens. Matter* **9**, 9021 (1997).
 - [32] H. Otsuka, *Phys. Rev. B* **53**, 14004 (1996).
 - [33] L. Tagliacozzo, G. Evenbly, and G. Vidal, *Phys. Rev. B* **80**, 235127 (2009).
 - [34] V. Murg, F. Verstraete, and J. I. Cirac, *Phys. Rev. A* **75**, 033605 (2007).
 - [35] V. Murg, F. Verstraete, and J. I. Cirac, *Phys. Rev. B* **79**, 195119 (2009).
 - [36] D. Muth, B. Schmidt, and M. Fleischhauer, *New J. Phys.* **12**, 083065 (2010).
 - [37] V. Murg, F. Verstraete, O. Legeza, and R. M. Noack, *Phys. Rev. B* **82**, 205105 (2010).
 - [38] S. Montangero, M. Rizzi, V. Giovannetti, and R. Fazio, *Phys. Rev. B* **80**, 113103 (2009).
 - [39] P. Corboz, G. Evenbly, F. Verstraete, and G. Vidal, *Phys. Rev. A* **81**, 010303 (2010).
 - [40] P. Corboz, R. Orús, B. Bauer, and G. Vidal, *Phys. Rev. B* **81**, 165104 (2010).
 - [41] P. Corboz, M. Lajkó, A. M. Läuchli, K. Penc, and F. Mila, *Phys. Rev. X* **2**, 041013 (2012).
 - [42] P. Corboz, S. R. White, G. Vidal, and M. Troyer, *Phys. Rev. B* **84**, 041108 (2011).
 - [43] S. Yan, D. A. Huse, and S. R. White, *Science* **332**, 1173 (2011).
 - [44] S. Depenbrock, I. P. McCulloch, and U. Schollwöck, *Phys. Rev. Lett.* **109**, 067201 (2012).
 - [45] G. Vidal, *Phys. Rev. Lett.* **99**, 220405 (2007).
 - [46] G. Vidal, *Phys. Rev. Lett.* **101**, 110501 (2008).
 - [47] G. Vidal, in *Understanding Quantum Phase Transitions*, edited by L. D. Carr (Taylor & Francis, Boca Raton, FL, 2010).
 - [48] G. Evenbly and G. Vidal, *J. Stat. Phys.* **145**, 891 (2011).
 - [49] B. Swingle, *Phys. Rev. D* **86**, 065007 (2012).
 - [50] B. Swingle, [arXiv:1209.3304v1](#) [hep-th].
 - [51] S. Singh and G. Vidal, *Phys. Rev. B* **88**, 121108 (2013).
 - [52] G. Evenbly and G. Vidal, *Phys. Rev. B* **89**, 235113 (2014).
 - [53] X. Chen, Z.-C. Gu, and X.-G. Wen, *Phys. Rev. B* **83**, 035107 (2011).
 - [54] X. Chen, Z.-C. Gu, and X.-G. Wen, *Phys. Rev. B* **84**, 235128 (2011).
 - [55] N. Schuch, D. Pérez-García, and I. Cirac, *Phys. Rev. B* **84**, 165139 (2011).
 - [56] L. Cincio and G. Vidal, *Phys. Rev. Lett.* **110**, 067208 (2013).
 - [57] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
 - [58] U. Schollwöck, *Ann. Phys.* **326**, 96 (2011).
 - [59] G. Vidal, *Phys. Rev. Lett.* **93**, 040502 (2004).
 - [60] G. Vidal, *Phys. Rev. Lett.* **98**, 070201 (2007).
 - [61] Y.-Y. Shi, L.-M. Duan, and G. Vidal, *Phys. Rev. A* **74**, 022320 (2006).
 - [62] F. Verstraete and J. I. Cirac, [arXiv:cond-mat/0407066v1](#).
 - [63] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **101**, 250602 (2008).
 - [64] G. Evenbly and G. Vidal, *Phys. Rev. Lett.* **112**, 240502 (2014).
 - [65] M. Aguado and G. Vidal, *Phys. Rev. Lett.* **100**, 070404 (2008).
 - [66] R. N. C. Pfeifer, [arXiv:1304.6112v1](#) [cond-mat.str-el].
 - [67] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevE.90.033315> for the reference implementation of the NETCON algorithm.
 - [68] G. Evenbly and R. N. C. Pfeifer, *Phys. Rev. B* **89**, 245118 (2014).
 - [69] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. (MIT Press and McGraw-Hill, New York, 2001), Chap. 15.2, pp. 331–338.
 - [70] T. Hu and M. Shing, *SIAM J. Comput.* **11**, 362 (1982).
 - [71] T. Hu and M. Shing, *SIAM J. Comput.* **13**, 228 (1984).
 - [72] R. N. C. Pfeifer, G. Evenbly, S. Singh, and G. Vidal, [arXiv:1402.0939](#) [physics.comp-ph].