

RESEARCH ARTICLE

Benchmarking treewidth as a practical component of tensor network simulations

Eugene F. Dumitrescu¹, Allison L. Fisher², Timothy D. Goodrich^{1,2*}, Travis S. Humble¹, Blair D. Sullivan², Andrew L. Wright²

1 Quantum Computing Institute, Oak Ridge National Laboratory, Oak Ridge, TN, United States of America, **2** Department of Computer Science, North Carolina State University, Raleigh, NC, United States of America

* tdgoodri@ncsu.edu



OPEN ACCESS

Citation: Dumitrescu EF, Fisher AL, Goodrich TD, Humble TS, Sullivan BD, Wright AL (2018) Benchmarking treewidth as a practical component of tensor network simulations. PLoS ONE 13(12): e0207827. <https://doi.org/10.1371/journal.pone.0207827>

Editor: Emanuele G. Dalla Torre, Bar-Ilan University, ISRAEL

Received: June 29, 2018

Accepted: November 5, 2018

Published: December 18, 2018

Copyright: This is an open access article, free of all copyright, and may be freely reproduced, distributed, transmitted, modified, built upon, or otherwise used by anyone for any lawful purpose. The work is made available under the [Creative Commons CC0](https://creativecommons.org/licenses/by/4.0/) public domain dedication.

Data Availability Statement: All data are available at <https://github.com/TheoryInPractice/ConSequences>. Additionally, these data are available at Zenodo using the following DOI: and the Zenodo DOI is [10.5281/zenodo.1490522](https://doi.org/10.5281/zenodo.1490522).

Funding: Eugene Dumitrescu acknowledges support from the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U.S. Department of Energy. Allison Fisher, Blair D. Sullivan and Andrew Wright were supported in part

Abstract

Tensor networks are powerful factorization techniques which reduce resource requirements for numerically simulating principal quantum many-body systems and algorithms. The computational complexity of a tensor network simulation depends on the tensor ranks and the order in which they are contracted. Unfortunately, computing optimal contraction sequences (orderings) in general is known to be a computationally difficult (NP-complete) task. In 2005, Markov and Shi showed that optimal contraction sequences correspond to optimal (minimum width) tree decompositions of a tensor network's line graph, relating the contraction sequence problem to a rich literature in structural graph theory. While treewidth-based methods have largely been ignored in favor of dataset-specific algorithms in the prior tensor networks literature, we demonstrate their practical relevance for problems arising from two distinct methods used in quantum simulation: multi-scale entanglement renormalization ansatz (MERA) datasets and quantum circuits generated by the quantum approximate optimization algorithm (QAOA). We exhibit multiple regimes where treewidth-based algorithms outperform domain-specific algorithms, while demonstrating that the optimal choice of algorithm has a complex dependence on the network density, expected contraction complexity, and user run time requirements. We further provide an open source software framework designed with an emphasis on accessibility and extendability, enabling replicable experimental evaluations and future exploration of competing methods by practitioners.

Introduction

Tensor network factorizations provide a framework for controlled approximation which exponentially reduces the memory required to simulate a variety of quantum many-body systems [1, 2] and circuits [3, 4]. These factorizations do so by representing targeted sub-sectors of the full (exponentially scaling) Hilbert space. The tensors comprising the factorization are placed on the vertices of a graph, one appropriate to the geometry under consideration, and are contracted along the edges as needed to compute physical observables [1].

by the Gordon & Betty Moore Foundation's Data-Driven Discovery Initiative through Grant GBMF4560 to Blair D. Sullivan, and the NC State Provost's Professional Experience Program. Timothy D. Goodrich was funded with Government support under and awarded by DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a. Travis Humble acknowledges support from the U.S. Department of Energy, Office of Science Advanced Scientific Computing Research and Early Career Research programs. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Since their early usage as the density matrix renormalization group description for gapped spin chains [1, 5], tensor networks have been adapted and reformulated to also describe 2D area-law states [6, 7], critical systems [2], lattice gauge theories [8, 9], AdS/CFT duality [10], and open quantum systems [11, 12]. In addition to describing a wide range of physical phenomena, satisfiability problems [5, 13] and quantum computing simulations [3, 4] can be formulated as tensor contraction problems. In the case of the latter, simulations of quantum error correcting codes are leading to important insights into fault tolerant quantum computation [14, 15]. Additionally, given the tremendous interest in validating increasing complex experimental quantum computations [16], a flurry of simulations have recently appeared in which the underlying tensor network graph emerges from the structure of the algorithm being employed [17–20].

The overall descriptive power and algorithmic computational complexity, as formalized by the *contraction complexity* of a tensor network [4], is determined by the tensor network construction's underlying graph structure. For some tensor network algorithms (e.g. the matrix product state formulation) the contraction complexity is fixed and well understood. However, the task of determining the contraction complexity in general, along with computing an optimal contraction sequence witnessing this complexity, is NP-complete [21].

Despite this daunting theoretical complexity, efficient methods exist in practice for obtaining both optimal and 'good enough' contraction sequences. Domain-specific approaches typically search the space of all possible sequences and apply heuristic pruning techniques to reduce the search space [19, 21, 22]. Effective algorithms in this area incorporate pruning rules proprietary to the target application's data (e.g., MERA networks [22]), which limits their broader applicability. Another standard technique involves transforming the tensor network into a line graph, then computing a perfect elimination ordering and its treewidth, which can be translated into a contraction sequence and complexity for the original network, respectively [4].

In practice, however, engineering issues prevent these methods from being applied effectively. Domain-specific approaches typically suffer from proprietary construction, each assuming a different representation of the tensor networks, and using different code languages, dependencies, and interfaces (often with little-to-no documentation). Additionally, these implementations are typically only tested on the data for which they were designed, providing no expectation for how they might perform and/or scale in different contexts. Treewidth-based approaches further suffer from the graph theory overhead needed to convert their (typical) output of tree decompositions into perfect elimination orderings for the line graph and then contraction sequences for the tensor network.

Our primary contribution to the literature is to provide an open source code framework (available at github.com/TheoryInPractice/ConSequences) for integrating all existing contraction sequence algorithms into a common interface designed for extendability and documented for accessibility; further, we have tabulated the performance of several leading contraction sequence algorithms. Our results provide quantum circuit simulation developers an extended benchmark for expected performance on circuits with varying structures and complexities.

We use container-based (Docker [23]) wrappers for each contraction sequence algorithm, completely removing code dependency issues, and provide Python-based utilities for converting various input/output formats into standardized internal formats for consistency. We demonstrate the utility of this software by reproducing two previous studies based on domain-specific algorithms, and extending them to include treewidth-based solvers in a broader set of experimental results. We find that modern treewidth solvers from the recent PACE 2017 coding challenge [24] are both faster and have more consistent run times than the domain-specific

algorithms. This speed increase allows us to study larger datasets in both experiments, and provide more competitive comparisons of a tensor network simulator [19] against Microsoft's LIQUI> Hilbert space simulator [25]. In particular, we show that contraction sequence algorithms are no longer the major bottleneck in tensor network simulations, and there is immediate value in work improving the scalability of downstream contraction code.

The paper is organized as follows. We begin with relevant definitions and an overview of related work in the Background, then describe the functionality of our code framework and considerations for use and extension in ConSequences: An Accessible, Extendable Framework. In the subsequent MERA Applications section, we reproduce a study by Pfeifer et al. [22], evaluating their algorithm `netcon` alongside two treewidth algorithms from the PACE 2017 challenge (`freetdi` [26] and `meiji-e` [27]) on a dataset including multi-scale entanglement renormalization ansatz (MERA) networks [2]. We extend this initial comparison on a larger corpus of MERA networks, pushing the limits of these exact contraction sequence solvers on a new benchmark. In the Applications with QAOA Circuit Simulation section, we reproduce a study by Fried et al. [19], evaluating another treewidth-based solver (`quickbb` [28]) against `freetdi` and `meiji-e` on quantum circuits formulated with Farhi et al.'s quantum approximate optimization algorithm (QAOA) for MAXCUT on r -regular graphs [29]. In addition to contraction sequence comparisons, we simulate the tensor network with qTorch [19], noting the correlation between simulation time and contraction complexity, and providing an updated comparison with Microsoft's LIQUI> simulator [25]. We conclude with a summary and directions for future work.

Background

For a graph G , we use $V(G)$ and $E(G)$ to denote the sets of vertices and edges, respectively, and use $G[X]$ to denote the subgraph induced by a set of vertices X . We say two vertices u, v are *adjacent* if $(u, v) \in E(G)$, and call the set of all vertices adjacent to v its *neighborhood* $N(v)$. The *degree* of v is $|N(v)|$, and a graph is r -regular if every vertex has degree r .

Formally, a *tensor network* is represented by a graph whose vertices correspond to tensors and edges denote tensor contractions over tensor indices. A *contraction* of two tensors corresponds to an *edge contraction* in the graph, where two vertices u, v with respective neighborhoods $N(u), N(v)$ are replaced with a single vertex uv with neighborhood $(N(u) \cup N(v)) \setminus \{u, v\}$.

In the remainder of this section we define contraction complexity and its relationship to notions from structural graph theory including treewidth, then outline the methods used to generate MERA and QAOA tensor networks, which are used as data in our experiments.

Contraction complexity and treewidth

Simulation of a tensor network requires its contraction down to a single tensor, and the network's structure imposes certain lower bounds on the information that must be kept in the network, fundamentally captured in the notion of contraction complexity:

Definition (Contraction Complexity (cc)). A contraction sequence is an ordering of a tensor network's edges, and the complexity of a contraction sequence S is the largest degree of a merged vertex created by contracting the network according to S . The contraction complexity (cc) of a tensor network is the minimum complexity over all possible contraction sequences.

The run time complexity for simulating a tensor network is $O(\chi^k)$, where χ is a refinement parameter for bond dimensions [22] and $k \geq cc$ is the cost of the contraction sequenced executed. Since better contraction sequences can yield exponentially-faster simulations, it is of practical interest to quickly find sequences with minimum cost. Unfortunately, computing the

contraction complexity optimally is an NP-hard [21] optimization problem with strong ties to structural graph theory under the guise of treewidth and elimination orderings.

Definition (Treewidth (tw)). A tree decomposition of a graph G is a tree T with a function f mapping nodes in T to bags (sets) of vertices from G , such that the following conditions hold:

1. All vertices are represented: $\bigcup_{t \in V(T)} f(t) = V(G)$.
2. All edges are represented: $\forall (u, v) \in E(G), \exists t \in V(T)$ s.t. $u, v \in f(t)$.
3. Graph vertices induce a (connected) subtree of T : if $w \in f(r) \cap f(s)$ for $r, s \in V(T)$, $w \in V(G)$, then $w \in f(t)$ for all t on the path from r to s in T .

The width of a tree decomposition is $\max_{t \in V(T)} |f(t)| - 1$, and the treewidth of a graph G , denoted $tw(G)$, is the minimum width over all valid tree decompositions of G .

Perhaps surprisingly, treewidth can be viewed as a vertex-centric formulation of the edge-centric contraction complexity, via a transformation of the underlying graph G to its line graph $L(G)$. The line graph is constructed with $V(L(G)) = E(G)$ and $E(L(G)) = \{(e_1, e_2) \mid e_1 \neq e_2 \in E(G)$ s.t. e_1, e_2 share a common endpoint $\}$. The treewidth of $L(G)$ then captures the same complexity:

Theorem (Markov and Shi [4]). The contraction complexity of a graph equals the treewidth of its line graph.

The relationship between treewidth and contraction sequences is perhaps easier seen through the characterization of treewidth using *elimination orderings*, which are permutations of the vertices. Given an elimination order $\pi = v_1, v_2, \dots, v_n$ of a graph G , the *fill-in graph* G_π is constructed by iterating over v_i from $i = 1$ to n and adding edges to make the neighbors of v_i in $G[v_i, \dots, v_n]$ a clique. A graph has treewidth at most k if and only if there exist an elimination ordering π so that each vertex has at most k higher numbered neighbors in G_π (see e.g., [30] for a proof). This condition naturally corresponds to the maximum degree of a tensor in the contraction sequence.

While this intuition provides an straightforward mapping from a contraction sequence to a tree decomposition, the other direction of Markov and Shi's proof shows how to convert an arbitrary tree decomposition into a contraction sequence with equal complexity. This non-trivial conversion allows treewidth solvers, whose native outputs are only tree decompositions, to be used directly as contraction sequence algorithms. To enable future work, we provide an implementation of this conversion as a modular subroutine in our post-processing utilities.

Rapid advances have been made in treewidth solvers in recent years, in large part to the Parameterized Algorithms and Computational Experiments (PACE) Challenge [24, 31]. Previous algorithms with practical implementations (such as `quickbb` [28]) are based on searching the space of elimination orderings and given the connection between contraction sequences and elimination orderings [4], share a strong resemblance to typical domain-specific algorithms [21, 22].

However, recent work in separator-based treewidth algorithms has begun to dominate modern benchmarks. The classic Arnborg, Corniel, and Proskurowski dynamic programming algorithm [32], reformulated as a positive-instance dynamic programming (PID) algorithm, has produced the winners of both the PACE 2016 (a Java implementation by Tamaki) and PACE 2017 (a C++ implementation by Larisch and Salfelder, `freetdi`) challenges [24, 26, 31]. The 2017 challenge also saw a better scaling implementation (`meiji-e` [27]) based off of a PID-reformulation of the improved dynamic programming algorithm by Bouchitté and Todinca [33].

MERA tensor networks

One class of tensor networks that we examine is the multi-scale entanglement renormalization ansatz (MERA). Given that contraction sequence algorithms only utilize the structure of the underlying graph in these networks, we restrict our presentation here to the important structural notions (visualized with a 1D binary MERA in Fig 1). We direct the interested reader to [34] for a rigorous description beyond the graph structure.

Fundamentally, MERA is a scheme for mapping a lattice of *operator sites* onto a coarser lattice. This mapping is expressed in terms of *coarsening layers*. The lattices on which MERA acts have an inherent dimension, which we denote d ; for simplicity we only consider examples in 1- and 2-dimensions in this paper. The most detailed lattice (\mathcal{L}_0) contains all sites for operators, and lattice \mathcal{L}_1 is produced after one level of coarsening. A coarsening level consists of a layer of unitaries followed by a layer of isometries. To disentangle the sites, in the MERAs we consider, *unitary* tensors take in 2^d wires (edges) and output 2^d wires for a lattice of dimension




Fig 1. (Top) A 1D binary MERA with a 16-site lattice and 3 levels of coarsening; three operator placements are highlighted (red, blue, green). (Bottom) Causal cones and final tensor networks for each of the three highlighted operators. Note that the tensor networks for the left-most (red) and right-most (green) operators are isomorphic to one another, but structurally distinct from the middle (blue) operator's network.

<https://doi.org/10.1371/journal.pone.0207827.g001>

d. For the coarsening layer, k : 1 *isometry* tensors take in k wires and output one wire. In total, going from lattice \mathcal{L}_i with s sites to \mathcal{L}_{i+1} requires $\frac{s}{2}$ unitary tensors, $\frac{s}{2}$ isometry tensors, and produces a new lattice with $\frac{s}{k}$ sites. This structure is then reflected for negative lattice levels, and a wire connects the top and bottom level interface tensors. Once this MERA graph is defined, operators are placed on lattice sites in \mathcal{L}_0 and the *causal cone* is computed by including the operators and any tensor that lies on an ascending (descending) path to the upper (lower) interface tensor (Fig 1). Once the causal cone is computed, every tensor not included in the cone is removed (by unitarity and properties of the isometries), and wires are added from a tensor to its dual mirror such that all tensors have the requisite number of wires.

QAOA quantum circuits

Another source for data comes from the quantum approximation optimization algorithm (QAOA) [29], a hybrid classical-quantum algorithm for utilizing near-term (~ 100 qubit) quantum computers. While applicable to generic satisfiability problems, we restrict ourselves to the MAXCUT optimization problem on r -regular graphs. Notably, when QAOA is applied to the MAXCUT problem, the structure of the input graph is reflected in the quantum circuit. Interested readers should refer to Farhi et al.'s formulation of QAOA [29] for a theoretical treatment, or the qTorch source code [19] for a practical example.

ConSequences: An accessible, extendable framework

One issue preventing widespread experimentation with (and adoption of) contraction sequence algorithms was the practical problem of installing the software and managing software dependencies. Of the algorithms presented in this paper, one is interfaced with MATLAB and uses C extensions for computationally-difficult sections (`netcon`), one is written in Java (`meiji-e`), two are written in C++ (`freetdi`, `qTorch`), and one is *only* distributed as a binary executable for Linux (`quickbb`). Often these implementations were written as a proof of concept and contain little-to-no documentation, especially regarding the code library dependencies needed to compile the code.

Additionally, once the code is compiled, each solver has proprietary input and output format. Algorithms from the treewidth literature may require the input graph to have particular vertex labels, and typically output a tree decomposition or an elimination ordering. Algorithms from the contraction sequence literature may require the input as a quantum circuit in Quantum Assembly format or as a tensor network, and the contraction sequence output may be an ordering of edges in the network or a sequence of contractions that automatically removes resulting self-loops.

Addressing both problems at once, we provide an open source framework `ConSequences` (Fig 2) for running contraction sequence algorithms, designed to be both *accessible* and *extendable*. The code is available at `github.com/TheoryInPractice/ConSequences`, complete with documentation for using existing code and tutorials for extending the functionality. At the core of this framework is a pipeline for pre-processing input, dispatching solvers, and post-processing output. The pre-processing utilities take tensor networks in formats such as Quantum Assembly and various graph formats, then generate graph files in a standardized format. The solver dispatcher manages calls to contraction sequence algorithms, and includes functionality such as parallelism, managing seeds, and handling timeouts. The post-processing utilities then take output from these solvers and generate a full set of output (including tree decompositions, perfect elimination orderings, and contraction sequences), allowing various aspects of the output to be reported and analyzed.

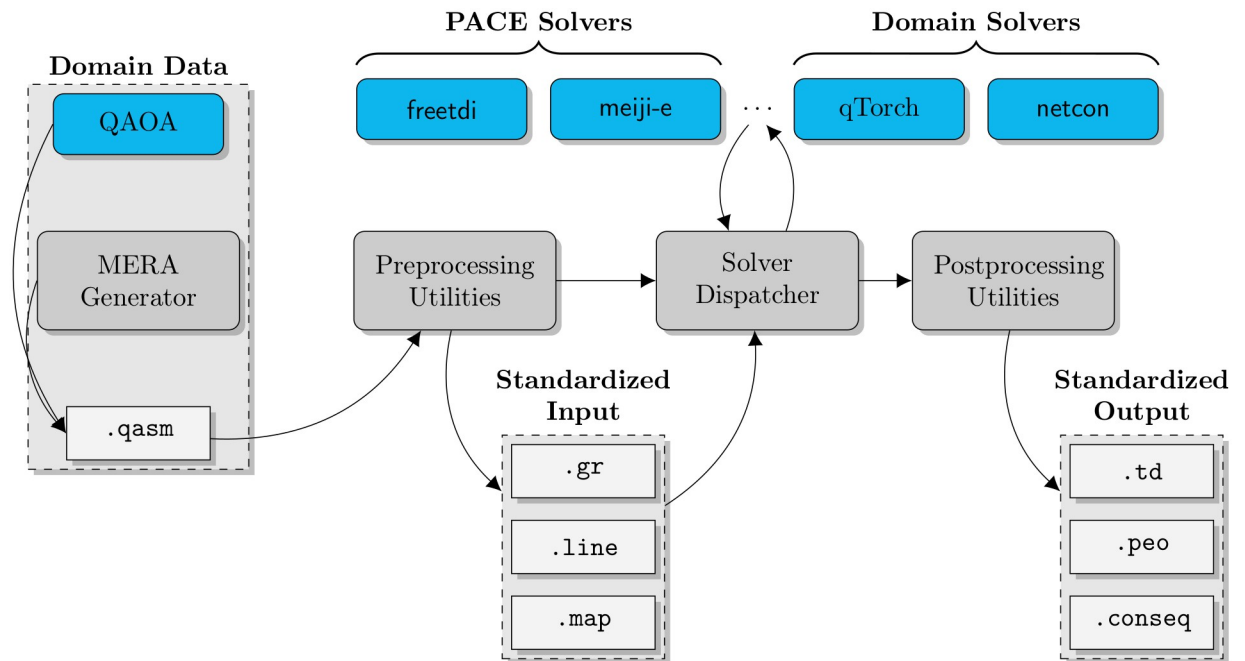


Fig 2. Visualization of the ConSequences pipeline. Externally-generated domain data is parsed into standardized graph formats with the pre-processing utility. The solver dispatcher then allows the user to compute contraction sequences (or their equivalent, e.g. tree decompositions) using external algorithms in Docker containers, then executes a post-processing utility to output standardized formulations of a contraction sequence.

<https://doi.org/10.1371/journal.pone.0207827.g002>

To make the framework *accessible*, the central pipeline is implemented as lightweight Python files, and individual contraction sequence algorithms are wrapped inside of Docker [23] images. Docker images are similar to virtual machines in that they abstract away all dependency issues, but have the distinct advantage of very little CPU and memory footprint on a native Linux system. Users need only install Python and Docker to run the code on Windows, MacOS or Linux; these steps are especially straightforward on Linux and we provide instructions for new users.

To make the framework *extendable*, we provide Docker templates for data generators (e.g., MERA networks) and contraction sequence algorithms. These templates are accompanied by tutorials which guide new users through wrapping up their code in Docker and interacting with our existing structure properly.

All experiments were run on three identical workstations, each with a single Xeon E5-2623 v3 processor (8 threads with a 3.0GHz base clock and 10MB cache) and 64GB system memory. Contraction sequence algorithms were run with a single thread and LIQUI|> and qTorch simulations with all threads. Experiments were run one at a time on the workstations, preventing noise from non-uniform cache usage between competing jobs. These workstations ran Fedora 27 with Docker 18.03.1-ce and Python 3.6.5. Algorithm-dependent software requirements (e.g., gcc, MATLAB) are fixed per algorithm in its Docker image wrapper; details are deferred to the code repository.

MERA applications

In this section we compare exact treewidth solvers to Pfeifer et al.'s *netcon* algorithm [22], on data from MERA networks. A small corpus of datasets from [22] is initially considered, in which case optimal contraction sequences are found within four seconds by both treewidth

solvers. To analyze algorithmic scaling with an extended benchmark, we generate 1- and 2-dimensional MERA networks with all possible placements of 1- and 2-operators, where we additionally observe that the `meiji-e` treewidth solver scales better than `freetdi` when networks become dense and increase in contraction complexity.

Initial comparison on `netcon` benchmark

The `netcon` implementation was developed as the contraction sequence algorithm for a simulation toolset written for MATLAB [22, 35, 36]. In the vein of previous approaches [21] such as depth-first search and dynamic programming, `netcon`'s core subroutine is a breadth-first search (BFS) over the solution space of all possible contraction sequences. To trim down this exponentially-sized space, the authors introduce two pruning methods for reducing the search space at each step of the BFS: first, if a contraction would cost more than a user-defined threshold, this contraction will not be considered; second, the authors provide a list of criteria for when outer product contractions should not be made. This algorithm is exact (i.e., it finds optimal contraction complexity), but its run time depends heuristically on the effectiveness of the pruning techniques to a particular network's search space. Provided as a MATLAB package, the core subroutines in `netcon` are implemented in external C code for efficiency. The authors of [22] evaluate their pruning heuristics on seven networks, including Tree Tensor Networks (TTN), Time-Evolution Block Decimation (TEBD), and MERA networks ranging from five to 27 tensors. In this previous work, the authors found that the pruning techniques resulted in faster results on all networks, with the largest network requiring 36 seconds.

We reproduce this experiment on all seven networks (provided in [22]). For `netcon` we use the MATLAB interface with an external C package, using all optimizations as specified in the accompanying code [22]. We compare against two exact PACE algorithms, `freetdi` and `meiji-e`. Refer to the ConSequences section for workstation specifications. As seen in Table 1, both `netcon` and `freetdi` require on the order of 0.001 seconds on the first four networks, whereas `meiji-e` was two orders of magnitude slower. On the three larger networks, however, `freetdi` was fastest, with both PACE algorithms finishing within four seconds on the largest network. This last data point is of particular interest because it hints at scalability differences between these algorithms. Whereas `meiji-e`'s run time was 5x slower on the 4:1 2D MERA compared to the 9:1 2D MERA, `freetdi` increased to over 250x slower and `netcon` over 450x slower.

Extended benchmark on large MERA networks

To test scalability further, we compared the algorithms' performance on a larger corpus of MERA networks. As seen in Fig 1, the iterative layers of isometries and unitaries in MERA

Table 1. Run times for each contraction sequence algorithm when executed on tensor network datasets from Pfeifer et al. [22]. For each tensor network, the number of tensors ($|V|$), edges ($|E|$), and optimal contraction complexity (cc) are reported.

Name	Instance			Run Time (sec)		
	$ V $	$ E $	cc	<code>freetdi</code>	<code>meiji-e</code>	<code>netcon</code>
3:1 1D TTN	5	28	6	0.009	0.338	0.002
TEBD	6	24	3	0.002	0.297	0.002
3:1 1D MERA	7	45	8	0.005	0.381	0.007
9:1 2D TTN	9	38	12	0.004	0.459	0.004
2:1 1D MERA	11	34	9	0.005	0.386	0.016
9:1 2D MERA	19	62	16	0.011	0.669	0.051
4:1 2D MERA	27	55	26	2.944	3.534	24.415

<https://doi.org/10.1371/journal.pone.0207827.t001>

networks allow one to easily generate underlying graphs given a unitary and isometry specification. We provide a MERA generator in our code for 1D lattices with binary isometries and 2D lattices with 4:1 isometries. This generator takes as input the number of coarsening levels and whether the top level of isometries should connect to a common tensor in the style of [34]. Once a MERA network is generated from these parameters, the locations of operators to be evaluated are chosen, a causal cone is computed, and the network is reduced down to the final tensor network by simplifying the network outside the causal cone.

Notably, given a fixed set of parameters needed to generate the MERA network, several choices of operator inputs can result in isomorphic (i.e., identical up to relabeling) networks. As seen in Fig 1, placing one operator at various sites can result in both isomorphic and non-isomorphic networks, so care must be taken to generate an appropriate set of representative graphs.

In our extended comparison, we generate 1D binary MERA and 2D 4:1 MERA with 1 and 2 operator placements. For 1 operator, we generate a MERA network for every possible operator placement, then compute the unique networks up to isomorphism. For 2 operators, we fix an operator at the first position (0 for 1D and (0,0) for 2D), then range the second operator over all possible positions; again we extract the unique graphs up to isomorphism. Fig 3 summarizes the number of networks up to isomorphism and details some of their nuances. For example, for a 1D 2:1 MERA with six coarsening levels and two operator placements, the network will have between 40 and 86 vertices based on operator placement (see the dark-blue highlighted row in the summary table). In this case, over all placements, $|\mathcal{M}| = 63$ networks were generated, but only $|\widehat{\mathcal{M}}| = 48$ were unique. Looking into these unique networks even further (left table), we find that only one network had 40 vertices, and over half of the unique graphs had over 80 vertices. While these networks may vary in the extreme cases, on average these networks are fairly predictable.

Figs 4 and 5 visualize the results of our extended experiment running contraction sequences algorithms on larger MERA networks using ConSequences. In Fig 4, networks are binned by the number of qubits they can support for a quantum simulation (which is determined by the number of possible operator sites) and run time is reported in seconds. For the networks that were included, a timeout of 20 minutes was used; if no algorithm could solve an instance within 20 minutes then the network was not included as a datapoint. From this perspective, `netcon` appears slower than both treewidth-based solvers, and `freetdi` generally dominates `meiji-e` in 1D MERAs. However, on 2D MERAs with one operator, all algorithms seem roughly equal for 16-qubits and 64-qubits, with `meiji-e` pulling slightly ahead on 256-qubits. The improved scaling of `meiji-e` is reflected in 2D MERAs with two operators, although here we begin to hit the limit of exact algorithms with a 20 minute timeout.

In Fig 5, networks are binned based on their contraction complexity, then run times are reported for each algorithm. From this perspective we find a stipulation for `freetdi`'s dominance on 1D MERAs: the contraction complexity never exceeded 9. Indeed, `freetdi` performed well on 2D MERAs until the optimal contraction complexity reaches 18, at which point `meiji-e` starts to outperform both algorithms. These experiments depict that finding the optimal contraction sequence within 20 minutes becomes problematic when the contraction complexity reaches the mid-twenties, which may be useful for predicting when heuristics should be used.

We note that while the experiment from [22] contained a graph with contraction complexity of 26, this network only had 27 tensors and 55 edges, whereas our 2D 4:1 MERA networks with contraction complexity 26 could have as many as 393 edges. This fact implies that search-space-based approaches need pruning techniques that prune exponentially-many sequences

1D 2:1 MERA				ℓ	$ V $	$ \mathcal{M} $	$ \widehat{\mathcal{M}} $
$\ell = 6, 2$ oper.				1D 2:1 MERA, 1 oper.			
$ V $	$ E $	$ \mathcal{S} $	$ \widehat{\mathcal{S}} $	2	15-17	4	2
40	65	1	1	3	21-27	8	3
62	104	2	1	4	27-37	16	8
64	107	2	1	5	33-47	32	21
66	110	2	1	6	39-57	64	50
68	113	2	1	7	45-67	128	111
70	116	2	1	8	51-77	256	236
70	117	2	1	9	57-87	512	489
72	120	2	1	10	63-97	1024	998
74	123	2	1	1D 2:1 MERA, 2 oper.			
74	124	4	3	2	16-22	3	2
76	127	4	3	3	22-38	7	4
78	130	4	3	4	28-54	15	9
80	134	8	7	5	34-70	31	21
82	137	8	7	6	40-86	63	48
86	144	16	15	7	46-102	127	106
				8	52-118	255	227
				9	58-134	511	475
				10	64-150	1023	978
				2D 4:1 MERA, 1 oper.			
				2	23-25	12	2
				3	29-55	68	7
				4	43-81	256	122
				2D 4:1 MERA, 2 oper.			
				2	29-38	9	3
				3	40-84	88	40
				4	81-132	223	134

2D 4:1 MERA			
$\ell = 4, 2$ oper.			
$ V $	$ E $	$ \mathcal{S} $	$ \widehat{\mathcal{S}} $
88	259	4	1
90	264	4	1
104	305	16	9
108	315	8	1
110	320	4	1
110	323	16	9
112	331	8	5
114	333	16	9
116	341	16	13
118	346	4	1
120	351	28	25
120	354	20	13
122	359	8	1
122	362	8	5
124	367	4	1
126	372	40	33
130	385	8	1
132	390	4	1
132	393	6	3

Fig 3. Summary of extended MERA benchmark data. (Center) For each lattice type (1D binary or 2D 4-ary) and number of operators, as the number of levels ℓ varies, we report the number of vertices $|V|$ in the resulting networks, the total number of networks $|\mathcal{M}|$, and the number of unique networks up to isomorphism $|\widehat{\mathcal{M}}|$. (Left, Right) These detailed tables each expand a row from the summary table, specifying the number of networks produced (total $|\mathcal{S}|$ and up to isomorphism $|\widehat{\mathcal{S}}|$), for each pair of values for the number of vertices $|V|$ and number of edges $|E|$. Note that sum of the $|\mathcal{S}|$'s in a detailed table sums to the corresponding $|\mathcal{M}|$ in the summary table, and likewise for $|\widehat{\mathcal{S}}|$ and $|\widehat{\mathcal{M}}|$.

<https://doi.org/10.1371/journal.pone.0207827.g003>

in the number of network edges, otherwise the run time will scale without the optimal contraction complexity necessarily increasing.

Applications with QAOA circuit simulation

In this section we compare exact treewidth solvers from PACE with the `quickbb` solver used in Fried et al.'s `qTorch` tensor network simulator [19]. This comparison is run on quantum circuits constructed with the QAOA method for solving `MAXCUT` on r -regular graphs. We find that exact contraction sequences on these graphs can be computed in less time than needed to execute the tensor contractions, allowing us to discuss the total simulation time without using an untimed preprocessing step. Rerunning the comparison with Microsoft's `LIQUI`>




Fig 4. Run times for the contraction sequence algorithms on select MERA networks, binned by number of qubits possible (number of \mathcal{L}_0 sites). All algorithms are timed out at 20 minutes (horizontal dashed line), and a network that remained unsolved by every algorithm is not included. 2D MERA with one operator had 48 of 131 networks that did not finish, and the two operator networks had 193 of 207 that did not finish.

<https://doi.org/10.1371/journal.pone.0207827.g004>

simulator from [19], we find that tensor networks are competitive against state-of-the-art simulators, allowing speedups when little information is needed to represent the network (i.e., on sparse graphs), and a potential solution for the 22-qubit limit currently built into LIQUi>.

Computing contraction complexity

Before comparing total simulation times, we start by evaluating the times required to find optimal contraction sequences. Because non-optimal contraction sequences lead to exponentially-slower downstream simulation times, we are particularly interested in exploring the limits of exact solvers.

The data for this experiment comes from qTorch’s QAOA quantum circuit constructor, which computes MAXCUT on a specified (arbitrary) graph. Reproducing circuits similar to the original qTorch experiments, we use r -regular graphs for $r \in \{3, 4, 5\}$, generated with NetworkX and seeded for easy replicability.

Fig 6 visualizes the results from using 25 random graphs for each $(r, |V|)$ pair and a 15 minute timeout. Similar to the sparse MERA graphs, we find that *freetdi* dominates run times and is up to five orders of magnitude faster than *quickbb* on the smallest graphs. *meiji-e* again scales better than *freetdi*, allowing it to compute optimal contraction sequences for several networks within the timeout that other algorithms could not finish. Notably, *quickbb* is an anytime algorithm that finds increasingly better perfect elimination orderings, so it always provides a (potentially non-optimal) solution when given a timeout. Similar behavior may be adapted from heuristic versions of PACE submissions (e.g., *meiji-e*), but this functionality is left as future work.




Fig 5. Run times for the three algorithms on select MERA networks, binned by optimal contraction complexity. All algorithms are timed out at 20 minutes (horizontal dashed line), and a network that remained unsolved by every algorithm is not included. 2D MERA with one operator had 48 of 131 networks that did not finish, and the two operator networks had 193 of 207 that did not finish.

<https://doi.org/10.1371/journal.pone.0207827.g005>

Table 2 overviews the optimal contraction complexity for r -regular graphs, $|V|$ vertices, 25 samples per $(r, |V|)$ pair. Optimal values are guaranteed and found by at least one of *freetdi* and *meiji-e*. We find these values have a small standard deviation, implying that the quantum circuit resulting from QAOA computing MAXCUT have fairly predictable contraction complexity, which may be useful for recognizing a priori the time and hardware resources needed to evaluate a tensor network.

Simulation run times

Computing downstream simulation times using qTorch, we first evaluate how simulation time correlates with contraction complexity. This comparison enables us to quantify how downstream runtime is impacted by the contraction complexity in practice (theoretical analysis predicts an exponential increase due to the size of the merged tensors, but it is possible that some downstream code mitigates this impact). Further, it provides for a finer-grained comparison of the quality of contraction sequences produced by *freetdi* and *meiji-e*, as measured by resulting simulation run time. Although both algorithms produce sequences with the same contraction complexity (and thus have the same leading-order term in the simulation time complexity), this term may occur during the contraction sequence between 1 and $|V|$ times. By examining raw simulation times, we may infer that one algorithm tends towards ‘higher quality’ contraction sequences than another.

As shown in Fig 7, simulation time indeed has an exponential dependence on the contraction complexity. Additionally, the contraction sequences produced by *freetdi* and *meiji-e* appear to be of comparable quality over the same corpus of networks. We observe

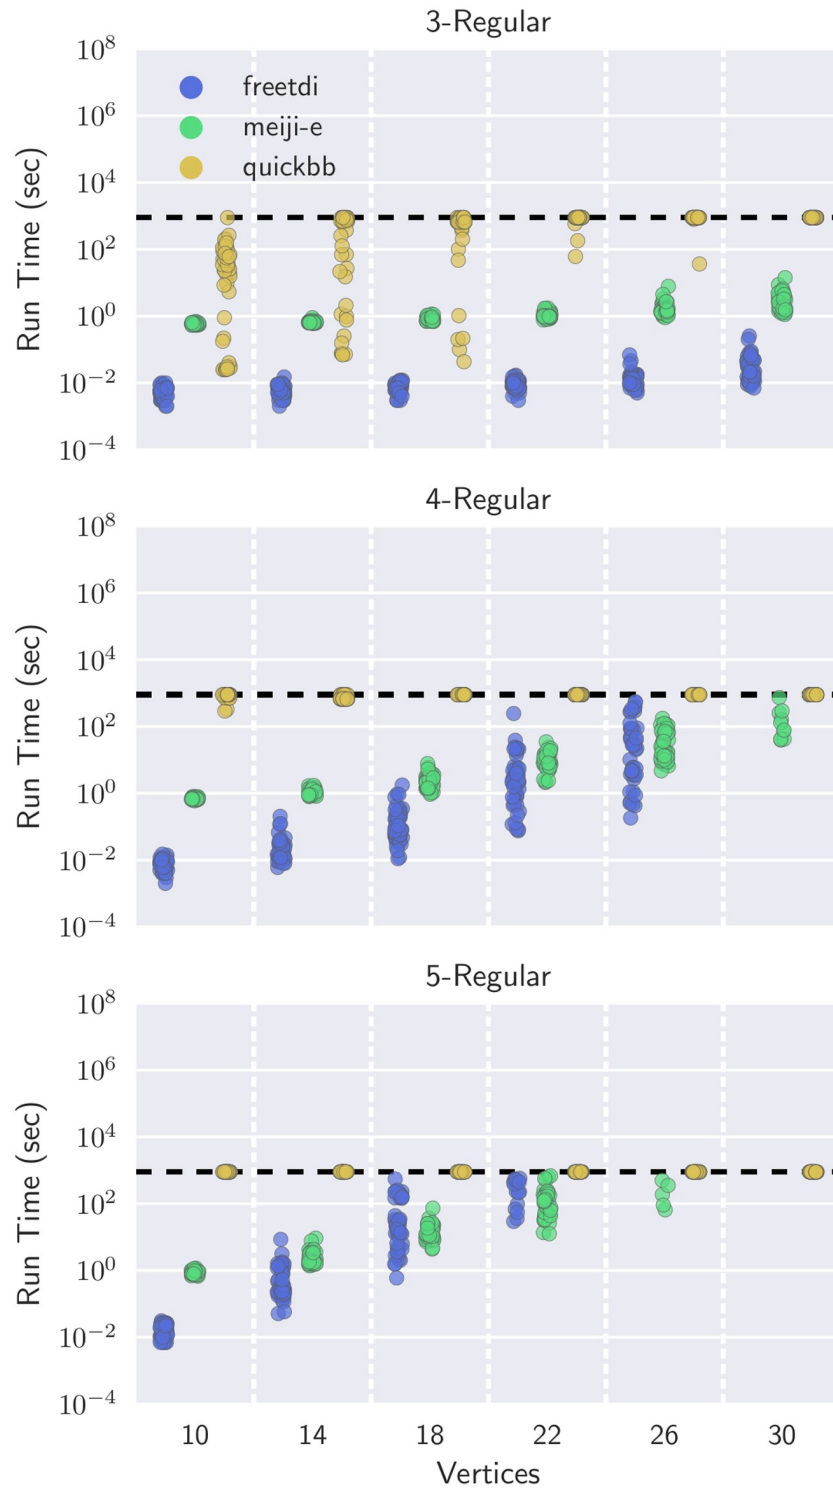


Fig 6. Run times for *freetdi*, *meiji-e*, and *quickbb* on QAOA circuits for computing MAXCUT on r -regular graphs. 25 random regular graphs are generated at each $r, |V|$ level using NetworkX, and algorithms were timed out at 15 minutes (horizontal line).

<https://doi.org/10.1371/journal.pone.0207827.g006>

Table 2. Exact contraction complexities found using `freetdi` and `meiji-e` on QAOA circuits for computing MAXCUT on r -regular graphs. 25 random regular graphs are generated at each $r, |V|$ level using NetworkX, and algorithms were timed out at 15 minutes. Timed out values were dropped from the data, resulting in less than 25 Samples for some parameter values.

Network		Optimal Contraction Complexity					
r	$ V $	Samples	Mean	S.D.	Min	50%	Max
3	10	25	5.0	0.6	4	5	6
	14	25	5.2	0.6	4	5	6
	18	25	6.0	0.8	5	6	7
	22	25	6.3	0.9	5	6	8
	26	25	6.7	0.7	6	7	8
	30	25	7.8	0.6	7	8	9
4	10	25	6.5	0.5	6	7	7
	14	25	7.7	0.8	6	8	9
	18	25	8.7	0.7	8	9	10
	22	25	10.2	0.8	8	10	11
	26	24	11.2	1.2	9	12	13
	30	5	12.0	0.7	11	12	13
5	10	25	7.7	0.6	6	8	9
	14	25	9.6	0.7	8	10	11
	18	25	11.3	0.7	10	11	13
	22	23	12.7	0.8	11	13	14
	26	1	12.0	NaN	12	12	12

<https://doi.org/10.1371/journal.pone.0207827.t002>

that the range of run times for a given contraction complexity is nearly always an order of magnitude, meaning that the variance in run time scales proportional to the total run time. This observation may be useful for reliably predicting simulation time based on a network’s known contraction complexity, which may be useful for optimizing future simulators.

In our final experiment, we compare qTorch paired with optimal contraction sequences to Microsoft’s LIQUi|> simulator. Refer to the ConSequences section for workstation specifications. Previous comparisons to LIQUi|> included non-optimal contraction sequences found by quickbb, which may have caused the downstream simulation to be exponentially more expensive. Additionally, the experiments in [19] were structured so that quickbb was run for 3000 seconds per network ahead of time, which resulted in impractical total run times for lower levels of regularity.

Results for this experiment are visualized in Fig 8. Aligning with previous conclusions in [19], both 3-regular and 4-regular graphs are more quickly simulated on tensor networks than LIQUi|>. We additionally find that tensor networks can scale beyond the 22-qubit limit imposed on LIQUi|> (supposedly for exponential system memory usage). Even with 5-regular graphs, qTorch remains competitive as an alternative simulator.

Comparing Fig 6 with Fig 8, it is also clear that computing the contraction complexity is no longer the primary bottleneck. Instead, efforts should now be directed to improving the contraction simulation times.

Conclusion

In summary, ConSequences provides an open source, extendable platform for comparing contraction sequence algorithms for tensor networks. By packaging conversion utilities with containerized solvers, we remove both the theoretical and engineering difficulties preventing practitioners from running any contraction sequence solver on any tensor network.




Fig 7. Simulation time is tightly correlated with the contraction complexity of a network. While exact algorithms *freetdi* and *meiji-e* may generate different tree decompositions and thus contraction sequences with the same treewidth, the differences have little impact on simulation times.

<https://doi.org/10.1371/journal.pone.0207827.g007>

Additionally, we demonstrate the framework’s applicability by reproducing and significantly extending several prior empirical evaluations. With MERA networks, we introduce a more extensive and difficult benchmark dataset which allows identification of solutions that will scale (e.g., *freetdi* on 1D MERA networks), subtler performance differences between algorithms (e.g., how *meiji-e* scales better for larger contraction complexities), and areas where new approaches are needed (e.g., 2D, 2-operator MERA, where 193 of 207 networks timed out). With qTorch on QAOA data, we were able to validate the exponential dependence of run time on optimal contraction complexity, and produce a total simulation time more representative of the full pipeline required for simulation. In doing so, we illuminate the urgent need for improved contraction times when the contraction sequence is known.

For contraction sequence algorithms, several avenues offer promising future work. With large MERA networks we found that exact solvers for optimal contraction sequences had prohibitively high run times, a difficulty which may require using non-exact heuristics. Several treewidth-based algorithms have heuristic formulations in a different track of the PACE 2017 challenge [24], and our comparison against domain-specific algorithms suggests that these PACE submissions would be the natural starting point for an investigation into heuristics. Additionally, the use of heuristics involves a trade-off between finding a sequence with smaller complexity and the additional search time, which is not a problem for exact solvers. Exploring this trade-off in practical applications may be of interest.

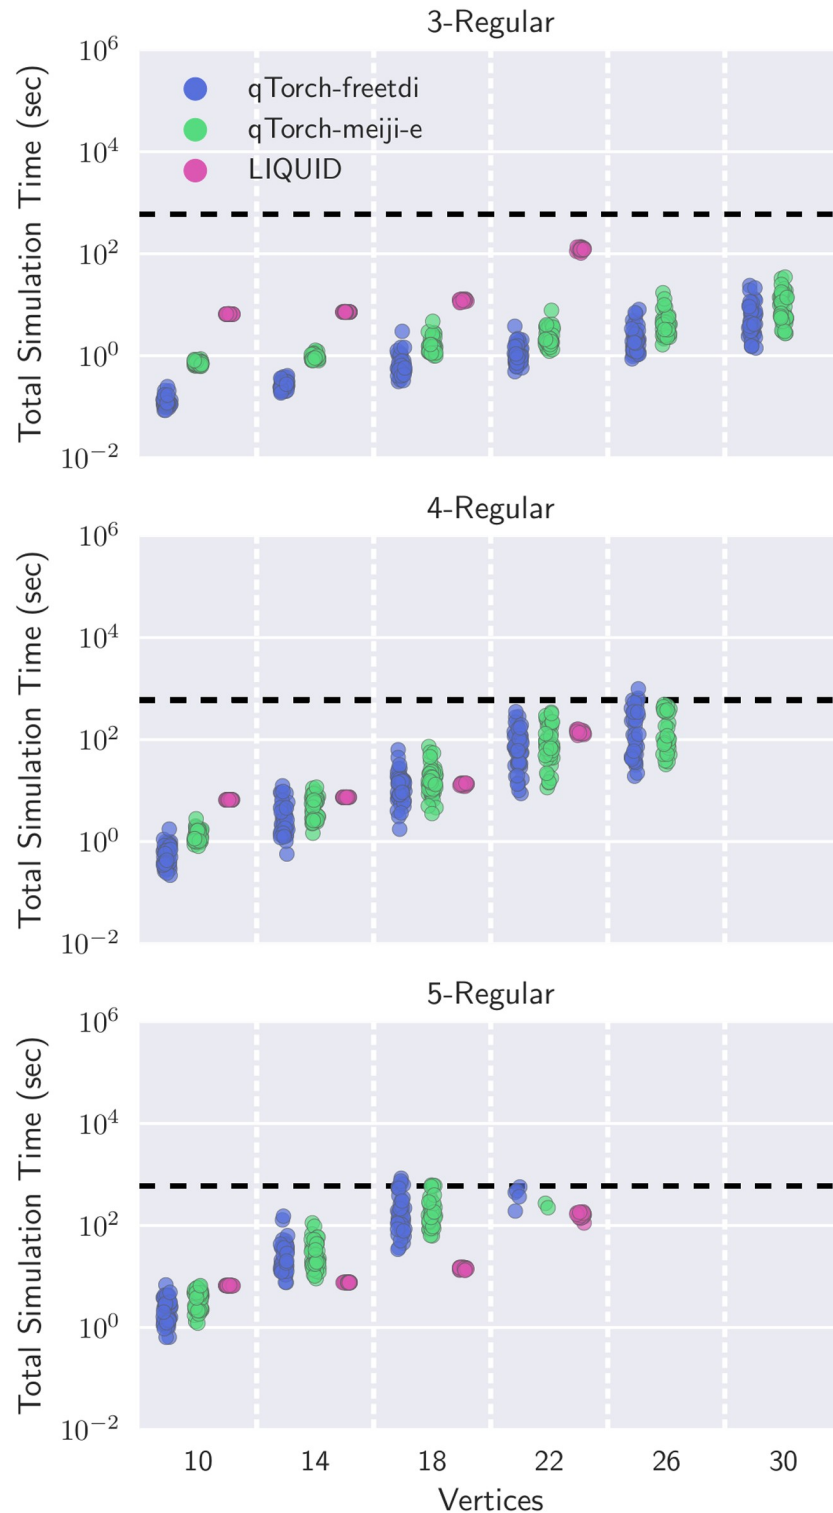


Fig 8. Simulation times of the qTorch tensor network simulator [19] with contraction sequences produced by exact treewidth algorithms vs. Microsoft’s LIQUID solver [25]. Total simulation time includes both computation of the contraction sequence using ConSequences and tensor network simulation time using qTorch. A timeout of 900 seconds is used for computing the contraction sequence (horizontal dashed line), and a simulation is not run unless an optimal contraction sequence is found by at least one contraction sequence algorithm. LIQUID is limited to simulations up to 22 qubits.

<https://doi.org/10.1371/journal.pone.0207827.g008>

Another consideration is that tensor network simulations are increasingly run on high-performance computing (HPC) systems. Modern HPC platforms scale-up performance with parallelism and heterogeneous computing accelerators (such as GPUs and FPGAs). As seen in recent work [37], adapting treewidth's state-of-the-art dynamic programming algorithms to these platforms is a non-trivial task, but would be of significant potential impact to the quantum computing community.

Acknowledgments

This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

Author Contributions

Conceptualization: Eugene F. Dumitrescu, Timothy D. Goodrich, Travis S. Humble, Blair D. Sullivan.

Data curation: Allison L. Fisher, Timothy D. Goodrich, Andrew L. Wright.

Funding acquisition: Eugene F. Dumitrescu, Timothy D. Goodrich, Travis S. Humble, Blair D. Sullivan.

Investigation: Allison L. Fisher, Timothy D. Goodrich, Blair D. Sullivan, Andrew L. Wright.

Methodology: Timothy D. Goodrich, Blair D. Sullivan.

Project administration: Blair D. Sullivan.

Software: Allison L. Fisher, Timothy D. Goodrich, Blair D. Sullivan, Andrew L. Wright.

Supervision: Blair D. Sullivan.

Validation: Allison L. Fisher, Timothy D. Goodrich, Andrew L. Wright.

Visualization: Allison L. Fisher, Timothy D. Goodrich, Blair D. Sullivan, Andrew L. Wright.

Writing – original draft: Allison L. Fisher, Timothy D. Goodrich, Blair D. Sullivan, Andrew L. Wright.

Writing – review & editing: Eugene F. Dumitrescu, Allison L. Fisher, Timothy D. Goodrich, Travis S. Humble, Blair D. Sullivan.

References

1. White SR. Density matrix formulation for quantum renormalization groups. *Physical Review Letters*. 1992; 69(19):2863–2866. <https://doi.org/10.1103/PhysRevLett.69.2863> PMID: 10046608
2. Vidal G. A class of quantum many-body states that can be efficiently simulated. *Phys Rev Lett*. 2006; 101(11):110501. <https://doi.org/10.1103/PhysRevLett.101.110501>
3. Vidal G. Efficient classical simulation of slightly entangled quantum computations. *Physical Review Letters*. 2003; 91(14):147902. <https://doi.org/10.1103/PhysRevLett.91.147902> PMID: 14611555
4. Markov IL, Shi Y. Simulating quantum computation by contracting tensor networks. *SIAM Journal on Computing*. 2005; 38(3):963–981. <https://doi.org/10.1137/050644756>

5. Biamonte J, Bergholm V. Tensor Networks in a Nutshell. arXiv preprint arXiv:170800006. 2017.
6. Eisert J, Cramer M, Plenio MB. Colloquium: Area laws for the entanglement entropy. *Rev Mod Phys*. 2010; 82:277–306. <https://doi.org/10.1103/RevModPhys.82.277>
7. Schuch N, Pérez-García D, Cirac I. Classifying quantum phases using matrix product states and projected entangled pair states. *Phys Rev B*. 2011; 84:165139. <https://doi.org/10.1103/PhysRevB.84.165139>
8. Rico E, Pichler T, Dalmonte M, Zoller P, Montangero S. Tensor Networks for Lattice Gauge Theories and Atomic Quantum Simulation. *Phys Rev Lett*. 2014; 112:201601. <https://doi.org/10.1103/PhysRevLett.112.201601>
9. Pichler T, Dalmonte M, Rico E, Zoller P, Montangero S. Real-Time Dynamics in U(1) Lattice Gauge Theories with Tensor Networks. *Phys Rev X*. 2016; 6:011023. <https://doi.org/10.1103/PhysRevX.6.011023>
10. Pastawski F, Yoshida B, Harlow D, Preskill J. Holographic quantum error-correcting codes: toy models for the bulk/boundary correspondence. *Journal of High Energy Physics*. 2015; 2015(6):149. [https://doi.org/10.1007/JHEP06\(2015\)149](https://doi.org/10.1007/JHEP06(2015)149)
11. Verstraete F, García-Ripoll JJ, Cirac JI. Matrix Product Density Operators: Simulation of Finite-Temperature and Dissipative Systems. *Phys Rev Lett*. 2004; 93:207204. <https://doi.org/10.1103/PhysRevLett.93.207204> PMID: 15600964
12. Werner AH, Jaschke D, Silvi P, Kliesch M, Calarco T, Eisert J, et al. Positive Tensor Network Approach for Simulating Open Quantum Many-Body Systems. *Phys Rev Lett*. 2016; 116:237201. <https://doi.org/10.1103/PhysRevLett.116.237201> PMID: 27341253
13. Biamonte JD, Morton J, Turner J. Tensor Network Contractions for #SAT. *Journal of Statistical Physics*. 2015; 160(5):1389–1404. <https://doi.org/10.1007/s10955-015-1276-z>
14. Ferris AJ, Poulin D. Tensor networks and quantum error correction. *Physical Review Letters*. 2014; 113(3):030501. <https://doi.org/10.1103/PhysRevLett.113.030501> PMID: 25083621
15. Darmawan AS, Poulin D. Tensor-Network Simulations of the Surface Code under Realistic Noise. *Physical Review Letters*. 2017; 119(4):40502. <https://doi.org/10.1103/PhysRevLett.119.040502>
16. Boixo S, Isakov SV, Smelyanskiy VN, Babbush R, Ding N, Jiang Z, et al. Characterizing quantum supremacy in near-term devices. *Nature Physics*. 2018; 14(6):1–6. <https://doi.org/10.1038/s41567-018-0124-x>
17. Dumitrescu E. Tree tensor network approach to simulating Shor’s algorithm. *Phys Rev A*. 2017; 96:062322. <https://doi.org/10.1103/PhysRevA.96.062322>
18. Dang A, Hill CD, Hollenberg LC. Optimising Matrix Product State Simulations of Shor’s Algorithm. arXiv preprint arXiv:171207311. 2017.
19. Fried ES, Sawaya NPD, Cao Y, Kivlichan ID, Romero J, Aspuru-Guzik A. qTorch: The Quantum Tensor Contraction Handler. arXiv preprint arXiv:170903636. 2017.
20. Pednault E, Gunnels JA, Nannicini G, Horesh L, Magerlein T, Solomonik E, et al. Breaking the 49-Qubit Barrier in the Simulation of Quantum Circuits. arXiv preprint arXiv:171005867. 2017.
21. Chi-Chung L, Sadayappan P, Wenger R. On optimizing a class of multi-dimensional loops with reduction for parallel execution. *Parallel Processing Letters*. 1997; 7(02):157–168. <https://doi.org/10.1142/S0129626497000176>
22. Pfeifer RNC, Haegeman J, Verstraete F. Faster identification of optimal contraction sequences for tensor networks. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*. 2014; 90(3):33315. <https://doi.org/10.1103/PhysRevE.90.033315>
23. Merkel D. Docker: lightweight linux containers for consistent development and deployment. *Linux Journal*. 2014; 2014(239):2.
24. Dell H, Komusiewicz C, Talmon N, Weller M. The PACE 2017 Parameterized Algorithms and Computational Experiments challenge: The second iteration. In: Lokshtanov D, Nishimura N, editors. *Leibniz International Proceedings in Informatics, LIPIcs*. vol. 89 of *Leibniz International Proceedings in Informatics (LIPIcs)*. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik; 2018. p. 30:1–30:12.
25. Wecker D, Svore KM. LIQUI>: A software design architecture and domain-specific language for quantum computing. arXiv preprint arXiv:14024467. 2014.
26. Larisch L, Salfelder F. FreeTDI PACE 2017 Submission; 2017. Available from: <https://github.com/freetdi/p17>.
27. Ohtsuka H, Tamaki H. Meiji PACE 2017 Submission; 2017. Available from: <https://github.com/TCS-Meiji/PACE2017-TrackA>.

28. Gogate V, Dechter R. A complete anytime algorithm for treewidth. In: Proceedings of the 20th conference on uncertainty in artificial intelligence. 1. AUAI Press; 2004. p. 201–208. Available from: <http://dl.acm.org/citation.cfm?id=1036868>.
29. Farhi E, Goldstone J, Gutmann S. A Quantum Approximate Optimization Algorithm. arXiv preprint arXiv:14114028. 2014.
30. Bodlaender HL. A partial k -arboretum of graphs with bounded treewidth. Theoretical Computer Science. 1998; 209(1):1–45. [https://doi.org/10.1016/S0304-3975\(97\)00228-4](https://doi.org/10.1016/S0304-3975(97)00228-4).
31. Dell H, Husfeldt T, Jansen BMP, Kaski P, Komusiewicz C, Rosamond FA. The first Parameterized Algorithms and Computational Experiments challenge. In: Guo J, Hermelin D, editors. Leibniz International Proceedings in Informatics, LIPIcs. vol. 63 of Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik; 2017. p. 30:1–30:9. Available from: <http://drops.dagstuhl.de/opus/volltexte/2017/6931>.
32. Arnborg S, Corneil DG, Proskurowski A. Complexity of finding embeddings in ak -tree. SIAM Journal on Algebraic Discrete Methods. 1987; 8(2):277–284. <https://doi.org/10.1137/0608024>
33. Bouchitté V, Todinca I. Treewidth and minimum fill-in: Grouping the minimal separators. SIAM Journal on Computing. 2001; 31(1):212–232. <https://doi.org/10.1137/S0097539799359683>
34. Evenbly G, Vidal G. Algorithms for entanglement renormalization ver 2. Physical Review B. 2009; 79(December 2008):1–17. <https://doi.org/10.1103/PhysRevB.79.144108>
35. Pfeifer RNC, Evenbly G, Singh S, Vidal G. NCON: A tensor network contractor for MATLAB. arXiv preprint arXiv:14020939. 2014.
36. Evenbly G, Pfeifer RNC. Improving the efficiency of variational tensor network algorithms. Physical Review B. 2013; 89(24):245118. <https://doi.org/10.1103/PhysRevB.89.245118>
37. Van Der Zanden TC, Bodlaender HL. Computing Treewidth on the GPU. arXiv preprint arXiv:170909990. 2017.