Tropical Tensor Network for Ground States of Spin Glasses

Jin-Guo Liu,[1,2,3,*] Lei Wang,[1,4,†] and Pan Zhang[5,6,7,‡]

[1]Beijing National Lab for Condensed Matter Physics and Institute of Physics, Chinese Academy of Sciences, Beijing 100190, China
[2]Harvard University, Cambridge, Massachusetts 02138, USA
[3]QuEra Computing Inc., Boston, Massachusetts 02143, USA
[4]Songshan Lake Materials Laboratory, Dongguan, Guangdong 523808, China
[5]CAS Key Laboratory for Theoretical Physics, Institute of Theoretical Physics, Chinese Academy of Sciences, Beijing 100190, China
[6]School of Fundamental Physics and Mathematical Sciences, Hangzhou Institute for Advanced Study, UCAS, Hangzhou 310024, China
[7]International Centre for Theoretical Physics Asia-Pacific, Beijing/Hangzhou, China

(Received 30 August 2020; revised 7 December 2020; accepted 26 January 2021; published 5 March 2021)

We present a unified exact tensor network approach to compute the ground state energy, identify the optimal configuration, and count the number of solutions for spin glasses. The method is based on tensor networks with the tropical algebra defined on the semiring of $(\mathbb{R} \cup \{-\infty\}, \oplus, \odot)$. Contracting the tropical tensor network gives the ground state energy; differentiating through the tensor network contraction gives the ground state configuration; mixing the tropical algebra and the ordinary algebra counts the ground state degeneracy. The approach brings together the concepts from graphical models, tensor networks, differentiable programming, and quantum circuit simulation, and easily utilizes the computational power of graphical processing units (GPUs). For applications, we compute the exact ground state energy of Ising spin glasses on square lattice up to 1024 spins, on cubic lattice up to 216 spins, and on three regular random graphs up to 220 spins, on a single GPU; we obtain exact ground state energy of $\pm J$ Ising spin glass on the chimera graph of D-Wave quantum annealer of 512 qubits in less than 100 s and investigate the exact value of the residual entropy of $\pm J$ spin glasses on the chimera graph; finally, we investigate ground-state energy and entropy of three-state Potts glasses on square lattices up to size 18×18. Our approach provides baselines and benchmarks for exact algorithms for spin glasses and combinatorial optimization problems, and for evaluating heuristic algorithms and mean-field theories.

DOI: 10.1103/PhysRevLett.126.090506

Introduction.—Combinatorial optimization problems are fundamental to theoretical studies in statistical physics and computer science. Efficient solutions to combinatorial optimization problems are also relevant to many practical applications such as operations research and artificial intelligence. A prototypical combinatorial optimization problem is finding the ground state of the Ising spin glass with the energy function

$$E(\{\sigma\}) = -\sum_{i<j} J_{ij}\sigma_i\sigma_j - \sum_i h_i\sigma_i, \tag{1}$$

where $\{\sigma\} \in \{\pm 1\}^N$ denotes a configuration of N Ising spins. Such a problem arises in broad contexts ranging from magnetic properties of dilute alloys [1] to probabilistic inference in graphical models [2]. Finding the ground state of the spin glass is non-deterministic polynomial-time (NP)-hard except on some special graphs [3]. This implies that an efficient solution to the problem is unlikely unless P = NP. Many NP problems have convenient Ising spin glass formulation [4]. In past decades, various approaches have been applied to such a problem, including simulated annealing on classical computers [5] and quantum annealing on manufactured quantum devices [6].

Besides the ground state energy and configurations, counting the number of ground-state configurations is also of interest from a physics and optimization perspective. The number of degeneracy characterizes the level of frustration and gives rise to residual entropy of the system at zero temperature [7]. For example, there can be an exponentially large number of degenerated ground states of the spin glass such that the system exhibits finite entropy density in the thermodynamic limit. Unfortunately, counting the number of the degenerated ground state of spin glasses is #P complete [8] which can be even harder than finding the ground state.

In this Letter, we present a unified approach to compute ground state energy, find out the ground state configuration, and count the ground state degeneracy of spin glasses exactly. The approach is based on the exact contraction of the tensor networks with tropical numbers which compute the spin-glass partition function directly in the zero-temperature limit. In principle, the approach is not conceptually new since there can be equivalent dynamic programming or message passing formulations. It is rather a synthesis of techniques in

combinatorial optimization, graphical models, and machine learning into a unified framework in the language of tensor networks, which provides valuable insights for efficient and generic implementations. In particular, the tropical tensor network offers a general computational framework so that one can easily exploit software and hardware advances in quantum circuit simulations, automatic differentiation, and hardware accelerations. In this regard, the approach adds another example along the fruitful line of research bridging graphical models, tensor networks, and quantum circuits [9–17].

There were previous efforts of investigating low-temperature properties of spin-glasses using approximated tensor contraction methods [17–19]. Among other things, these approaches and the related transfer matrix approach [20,21] face numerical issues at low temperatures due to the cancellation of tensor elements with exponential scales [22]. References [23,24] investigated the residual entropy of infinite translational invariant frustrated classical spin systems by constructing tensor networks according to local rules of the ground-state manifold. More closely related to the present Letter, one can employ exact tensor network contraction to count the number of solutions in the constraint satisfaction problems [25–28], however, with the ground-state energy known to be zero *a priori*.

Tropical tensor network.—Tropical algebra is defined by replacing the usual sum and product operators for ordinary real numbers with the max and sum operators, respectively, [29]

$$x \oplus y = \max(x, y), \qquad x \odot y = x + y. \tag{2}$$

One sees that $-\infty$ acts as the zero element for the tropical number since $-\infty \oplus x = x$ and $-\infty \odot x = -\infty$. On the other hand, 0 acts as the multiplicative identity since $0 \odot x = x$. The $\oplus$ and $\odot$ operators still have commutative, associative, and distributive properties. However, since there is no additive inverse, the $\oplus$ and $\odot$ and operations define a semiring over $\mathbb{R} \cup \{-\infty\}$. The semiring formulation unifies a large number of inference algorithms in the graphical models based on dynamic programming [30,31]. Recently, there have been efforts in combing the semiring algebra with modern deep learning frameworks with optimized tensor operations and automatic differentiation [32,33].

One can consider tensor networks whose elements are tropical numbers with the algebra Eq. (2). Since the elementary operations involved in contracting tensor networks are just sum and product, the contraction of tropical tensor networks is well defined. One can use such contraction to solve the ground state of the Ising spin glass. For example, consider the Ising spin glasses Eq. (1) defined on a two-dimensional square lattice, the tropical tensor network is shown in Fig. 1(a). The tensor network representation corresponds to the factor graph of the spin-glass

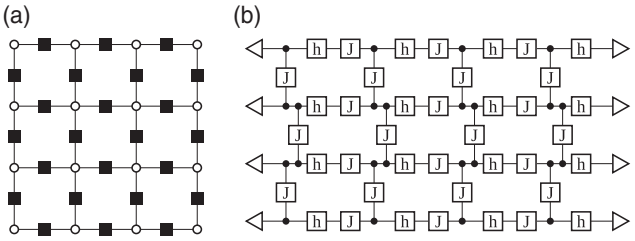

FIG. 1. (a) The tensor network representation of a square lattice Ising spin glass. (b) An equivalent circuit representation used for the practical simulation. See text for the definition of the symbols.

graphical model [30]. There are 2×2 tropical tensors $\blacksquare = \begin{pmatrix} J_{ij} & -J_{ij} \\ -J_{ij} & J_{ij} \end{pmatrix}$ that reside on the bond connecting vertices i and j, with the tensor elements being the negative coupling energies. The dots are diagonal tensors with $1 \circ 1 = h_i$, $2 \circ 2 = -h_i$, and $-\infty$ for all other tensor elements. In cases where the local field vanishes, these dots reduce to the copy tensor in terms of the tropical algebra which demands that all the legs have the same indices. Contraction of the tensor network under the tropical algebra gives the ground state energy of the Ising spin glass. In the contraction, the $\oplus$ operator selects the optimal spin configuration, and the $\odot$ operator sums the energy contribution from subregions of the graph. The intermediate tensors record the minimal energy given the external tensor indices, so they correspond to max marginals in the graphical model [34].

From a physics perspective, the tropical tensor network naturally arises from computing the zero-temperature limit of the partition function $Z = \sum_{\{\sigma\}} e^{-\beta E}$. The ground state energy, $E^* = -\lim_{\beta\to\infty}(1/\beta)\ln Z = -\lim_{\beta\to\infty}(1/\beta)\times \ln\sum_{\{\sigma\}}\prod_{i<j} e^{\beta J_{ij}\sigma_i\sigma_j}\prod_i e^{\beta h_i\sigma_i}$, involves ordinary sum and product operations for the Boltzmann weights. When taking the zero temperature limit, it is more convenient to deal with the exponents directly,

$$\lim_{\beta\to\infty}\frac{1}{\beta}\ln(e^{\beta x} + e^{\beta y}) = x \oplus y, \qquad \frac{1}{\beta}\ln(e^{\beta x} \times e^{\beta y}) = x \odot y, \tag{3}$$

which leads to the tropical algebra Eq. (2). The tropical representation also corresponds to the logarithmic number system [35] which avoids the numerical issue in dealing with exponentially large numbers on computers with finite precision numerics [22].

Moreover, one can also employ the present approach to count the number of ground states at the same computational complexity of computing the ground state energy. To implement this, we further generalize the tensor element to be a tuple (x, n) composed by a tropical number x and an

ordinary number n. The tropical number records the negative energy, while the ordinary number counts the number of minimal energy configurations. For tensor network contraction, we need the multiplication and addition of the tuple: $(x_1,n_1)\odot(x_2,n_2)=(x_1+x_2,n_1\times n_2)$ and

$$(x_1,n_1)\oplus(x_2,n_2)=\begin{cases}(x_1\oplus x_2,n_1+n_2) & \text{if } x_1=x_2\\ (x_1\oplus x_2,n_1) & \text{if } x_1>x_2\\ (x_1\oplus x_2,n_2) & \text{if } x_1<x_2\,.\end{cases} \tag{4}$$

Essentially, these two numbers in the tuple correspond to the leading order and the $\mathcal{O}(1/\beta)$ contributions (energy and entropy) in the low-temperature expansion of the log-partition function. After contracting the tensor network, one reads out the ground state energy and degeneracy from the two elements of the tuple. In this way, one can count the number of optimal solutions exactly without explicitly enumerating the solutions [36,37].

Contract tropical tensor networks.—We have formulated the computation of the ground state energy and the ground state degeneracy of the Ising spin glass Eq. (1) as a contraction of the tropical tensor network. On a tree graph, contraction of the tropical tensor network is equivalent to the max-sum algorithm [2], i.e., the maximum of a posterior version of the sum-product (belief propagation) algorithm on graphical models. On a general graph, when the junction tree algorithm [38] applies it can be treated as a special case of the tropical tensor network contraction algorithm using a specific contraction order utilizing a tree decomposition of the graph.

The contraction of a general tensor network belongs to the class of #P hard problems [39], so it is unlikely to find polynomial algorithms for exact contractions. Algorithmically, the computational complexity of tensor network contraction is exponential to the tree width of the network [9]. On a regular graph (e.g., 2D lattice), one can easily find a good contraction order that has an optimal computational complexity. However, on a general graph, a good contraction order is usually difficult to find, thus one usually relies on heuristic algorithms to identify a contraction order with low computational complexity. Reference [9] proposed to use tree decomposition of the line graph of the tensor network, found by a branch and bound algorithm. This has been widely adopted in subsequent works on classical simulation of quantum circuits with tensor networks [14,40–46]. Recently, more advanced heuristic algorithms have been developed by combining graph partition algorithms and greedy algorithms [47,48].

In addition to a good contraction order, efficient linear algebra libraries are also important for the performance of the contractions. For ordinary contractions, the basic linear algebra subprograms (BLAS) library is a standard tool for performing efficient product and plus operations, and can fully release the computational power of specialized hardware such as graphical processing units (GPUs) and tensor processing units. For tropical algebra, fortunately, basic operations can be inherited from standard linear algebra libraries as long as they are programmed in a generic manner to support $\oplus$ and $\odot$ operators. When performing contractions on GPUs, another important factor is memory efficiency, that is, all operations should be performed in-place without allocating extra memory. This actually shares the same demand as the simulation of quantum circuits. To this end, one can actually contract tropical tensor networks by repurposing software that was originally developed for quantum circuit simulations.

To sum up, the tropical tensor network formulation opens a way to leverage recent algorithmic and software advances in tensor network contraction for combinatorial optimization problems. Moreover, the tensor contraction formation fits nicely to the specialized hardware such as GPUs, where, as we reported below, one can actually employ low precision floating numbers (or even integer type for integral couplings) for better numerical performance and reduced memory usage.

Obtaining ground states with automatic differentiation.—Given the way to compute the ground state energy of the spin glass, there are several ways to obtain the ground state configurations. The most straightforward way would be running the same energy minimization program repeatedly with perturbed fields. Since the ground state energy is a piecewise linear function of the fields, the numerical finite difference of the energy with respect to fields suffices to determine the ground state configurations [49]. Alternatively, one can impose an arbitrary order of the spin variables and compute the conditional probability of a variable being in the ground state given the previous ones, then sample the ground state configurations according to the conditional probability [34]. Both methods need to re-run the contraction algorithm $\mathcal{O}(N)$ times with the same memory cost as finding the ground state energy. One can nevertheless trade memory for computation time by caching intermediate contraction results and backtracking the computation for minimal energy configuration.

We employ the differentiable programming technique to differentiate through the tropical tensor network contraction [50]. To this end, we program the whole tensor network contraction in a differentiable way and compute the gradient of the contraction outcome with respect to the tensor elements using automatic differentiation. We note that the general idea of differentiating through a combinatorial optimization solver applies to cases beyond tropical tensor network contraction [51]. It is well known that there is a time-space trade-off in different ways of performing the automatic differentiation to a computer program [52]. The forward mode automatic differentiation (such as

ForwardDiff.jl [53]) has the same time and memory cost as the finite difference approach. While in the other extreme limit, the reverse mode automatic differentiation (such as Nilang.jl [54]) displays the $\mathcal{O}(1)$ computation overhead compared to the forward tensor contraction, and $\mathcal{O}(N)$ memory overhead. The time versus memory trade-off can be further controlled flexibly by using the checkpointing technique [52].

Applications.—We first apply the tropical tensor network approach to the Ising spin glasses on $L \times L$ square lattices, with the tensor network shown in Fig. 1(a). Interestingly, the computation of tensor network contraction is similar to evolving a quantum state under the action of local quantum gates, with the crucial difference that we are now dealing with nonunitary gates with the tropical algebra.

As shown in Fig. 1(b), the tensor network is cast into the expectation of a tropical circuit on the state vector of 2^L dimension. We denote $\langle\!\!-$ $= \begin{pmatrix} 0 \\ 0 \end{pmatrix}$ so that the initial and final states are both product state $\begin{pmatrix} 0 \\ 0 \end{pmatrix}^{\otimes L}$. The square symbols represent tropical gates, in which $-[\mathrm{J}]-$ $= \begin{pmatrix} J_{ij} & -J_{ij} \\ -J_{ij} & J_{ij} \end{pmatrix}$ and $-[\mathrm{h}]-$ $= \begin{pmatrix} h_i & -\infty \\ -\infty & -h_i \end{pmatrix}$ are single-site gates. The symbol a c / b d $[\mathrm{J}]$ denotes two site gates acting on neighboring sites. In fact, it is a diagonal tropical matrix $\mathrm{diag}(J_{ij}, -J_{ij}, -J_{ij}, J_{ij})_{ab,cd}$, with the off-diagonal elements set to $-\infty$. The order of operation of these diagonal gates to the state vector can be arbitrary. Exploiting this intimate connection, we employ the quantum programming software Yao.jl [55] to contract these tropical tensor networks [56]. It enables us to obtain the ground state energy of 1024 spins with external fields in about 590 s on a single Nvidia V100 GPU, with single-precision floating numbers Float32 for the tensor elements.

Next, we consider spin glass instances with $\pm J$ coupling and no external field on the chimera graph of the actual D-Wave device [6] shown in Fig. 2(a). The chimera graph consists of unit cells arranged in a square grid of the size of $L \times L$. Each unit cell contains eight spins forming a complete bipartite graph. Each group of four spins within the unit cell connects horizontally or vertically to the spins in the neighboring unit cells. We transform the chimera graph into a tensor network shown in Fig. 2(b) by exploiting its specific structure [57]. The red and blue circles are tropical copy tensors that represent a group of four Ising spins within each unit cell. The black tensor describes the intra-unit-cell couplings. While the red and blue squares denote the intercell interaction in the vertical and horizontal direction, respectively. These tensors are all 16×16 tropical matrices that contain the couplings between the original Ising spins. Such a tensor network

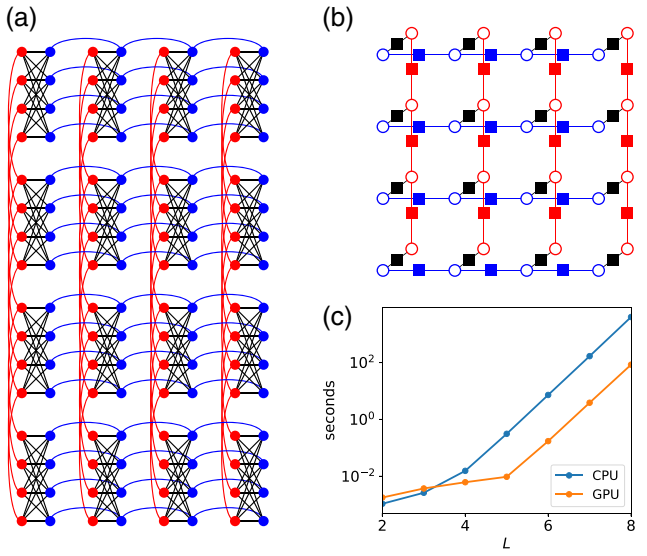

FIG. 2. (a) A chimera lattice with 4×4 unit cells. Dots represent Ising spins and lines indicate couplings. (b) Tensor network representation, where each node has a degree of freedom of four spins. (c) Wall clock time for computing the ground state energy of Ising spin glass on the chimera graph with the $L \times L$ unit cell ($8L^2$ spins).

formulation makes better use of the bipartite structure of the chimera graph than simply grouping the eight spins within the unit cell together [19]. After turning these tensors into local tropical gates, contraction of the tensor network can be carried out as evolution of a state with dimension 16^L. As shown in Fig. 2(c) one can obtain the ground state energy of $8L^2 = 512$ Ising spins in 84 s on the Nvidia V100 GPU. This is much faster than brute force enumeration using GPUs [58]. It is also slightly faster than the belief propagation exact solver running on 16 CPU cores used in Ref. [59]. We use Int16 data type for computational and memory efficiency, which is sufficient for such calculation since the energy has bounded integral values.

Figure 3(a) shows the histogram of the ground state degeneracy of the chimera spin glasses. One observes that the distributions are unimodal and broaden as the system size enlarges. Figure 3(b) shows the residual entropy density $s = \mathbb{E}[\ln g]/(8L^2)$ where g is the degeneracy and

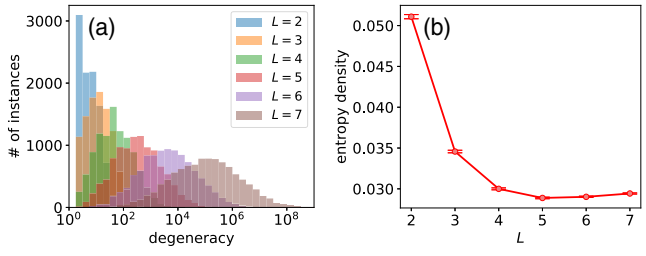

FIG. 3. (a) Histogram of the ground state degeneracy of $\pm J$ spin glasses on the chimera graph with $L \times L$ unit cells ($8L^2$ Ising spins). For each system size, we solve 10000 random instances. (b) The residual entropy density versus system size.

the expectation is over the 10000 random instances. The value of the residual entropy approaches $s = 0.03$ for increasingly larger system sizes. As a comparison, this value of the entropy density is smaller than the one of the $\pm J$ square lattice Ising spin glass $s \approx 0.07$ [21,60–64], indicating a smaller number of degenerated ground states on the chimera graph compared to the $\pm J$ square lattice spin glasses, possibly due to the larger connectivity in the Chimera graph which induces more constraints to each spin in the ground state and suppresses the degeneracy.

For problems on more general graphs, our method benefits from the contraction order developed in the quantum computation community [27,47,48,65,66]. As an example, with the present approach one can compute optimal solutions and count the number of solutions for spin glasses and combinatorial optimization problems on random graphs with hundreds of nodes, and check numerically the replica symmetry mean-field solutions [67,68]. Details can be found at Ref. [69].

Discussions.—An immediate implication of our method is that quantum circuit simulators can be repurposed to solve combinatorial optimization problems. This connection adds a profitable motivation for crafting efficient and generic quantum circuit simulators besides validating quantum devices.

We notice that the state-of-the-art method branch-and-cut approaches are able to achieve better performance for spin glasses on 2D lattices. For example, Ref. [70] reached 100×100 lattices for a spin glass with Gaussian couplings, and 50×50 lattices with $\pm J$ couplings [71]. However, the branch and bound method is less efficient in computing degeneracies. For example, the branch-and-bound results for entropy were reported with for 8×8 lattices [72], while, our method works out the ground-state entropy of $\pm J$ spin glass on 32×32 lattices. Moreover, the linear programming bounding method is sensitive to coupling types and connectivity of the model. On 2D lattices, the branch-and-cut method is quite efficient when equipped with the *circle inequality* [70] technique, especially with Gaussian couplings. But it turns out to be less efficient when the topology is a 3D lattice, where only results with $4 \times 4 \times 4 = 64$ spins are reported in the literature [72]. In contrast, on 3D lattices, our method works to $6 \times 6 \times 6 = 216$ spins. More seriously, if the model changes from an Ising spin glass to a Potts glass, not only the cutting plane method but also the linear programming bounding method breaks down. As a relief, one has to develop a more sophisticated semi-definite programming (SDP) method for providing energy lower bounds [73,74]. Reference [73] computed the ground-state energy of a $\pm J$ three-state Potts glass model on a 9×9 lattice using 10 h. As a comparison, our method is able to compute both ground-state energy and entropy on 18×18 lattices in 10 min, thus is significantly superior to SDP based branch-and-cut methods for Potts models [69]. Moreover, one could also apply specific bounds on the ground-state energy to enforce sparsity of the tropical tensors, this would combine the tropical tensor network framework with the branch and bound methods.

Moving forward, approximated contraction schemes for the tropical tensor networks may provide practical algorithms for the optimization and counting of large instances. A Julia implementation of the tropical tensor network used in this Letter is available at Ref. [75]. Thanks to generic programming, a minimalist working example contains only ~60 lines of code.

We thank Hai-Jun Liao, Zhi-Yuan Xie, and the BFS Tensor community for inspiring discussions, and Yingbo Ma for discussions on the Tropical BLAS library [76]. P. Z. is supported by project QYZDB-SSW-SYS032 of CAS, and Projects 12047503 and 11975294 of NSFC. L. W. is supported by the National Natural Science Foundation of China under Grant No. 11774398, and the Ministry of Science and Technology of China under Grants No. 2016YFA0300603 and No. 2016YFA0302400.

*cacate0129@gmail.com
†wanglei@iphy.ac.cn
‡panzhang@itp.ac.cn

[1] S. F. Edwards and P. W. Anderson, Theory of spin glasses, J. Phys. F **5**, 965 (1975).
[2] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques* (MIT Press, Cambridge, MA, 2009).
[3] F. Barahona, On the computational complexity of Ising spin glass models, J. Phys. A. **15**, 3241 (1982).
[4] A. Lucas, Ising formulations of many NP problems, Front. Phys. **2**, 5 (2014).
[5] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, Science **220**, 671 (1983).
[6] M. W. Johnson *et al.*, Quantum annealing with manufactured spins, Nature (London) **473**, 194 (2011).
[7] L. Pauling, The structure and entropy of ice and of other crystals with some randomness of atomic arrangement, J. Am. Chem. Soc. **57**, 2680 (1935).
[8] L. G. Valiant, The complexity of enumeration and reliability problems, SIAM J. Comput. **8**, 410 (1979).
[9] I. Markov and Y. Shi, Simulating quantum computation by contracting tensor networks, SIAM J. Comput. **38**, 963 (2008).
[10] A. Critch and J. Morton, Algebraic geometry of matrix product states, Symmetry Integr. Geom. Methods Appl. **10**, 095 (2014).
[11] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, Equivalence of restricted Boltzmann machines and tensor network states, Phys. Rev. B **97**, 085104 (2018).
[12] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, Unsupervised Generative Modeling Using Matrix Product States, Phys. Rev. X **8**, 031012 (2018).
[13] I. Glasser, N. Pancotti, and J. Ignacio Cirac, From probabilistic graphical models to generalized tensor networks for supervised learning, IEEE Access **8**, 68169 (2020).

[14] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, and H. Neven, Simulation of low-depth quantum circuits as complex undirected graphical models, arXiv:1712.05384.
[15] X. Gao, Z. Y. Zhang, and L. M. Duan, A quantum machine learning algorithm based on generative models, Sci. Adv. **4**, eaat9004 (2018).
[16] E. Robeva and A. Seigal, Duality of graphical models and tensor networks, Inf. Inference **8**, 273 (2019).
[17] F. Pan, P. Zhou, S. Li, and P. Zhang, Contracting Arbitrary Tensor Networks: General Approximate Algorithm and Applications in Graphical Models and Quantum Circuit Simulations, Phys. Rev. Lett. **125**, 060503 (2020).
[18] C. Wang, S. M. Qin, and H. J. Zhou, Topologically invariant tensor renormalization group method for the Edwards-Anderson spin glasses model, Phys. Rev. B **90**, 174201 (2014).
[19] M. M. Rams, M. Mohseni, and B. Gardas, Heuristic optimization and sampling with tensor networks for quasi-2D spin glass problems, arXiv:1811.06518.
[20] I. Morgenstern and K. Binder, Evidence against Spin-Glass Order in the Two-Dimensional Random-Bond Ising Model, Phys. Rev. Lett. **43**, 1615 (1979).
[21] H.-F. Cheung and W. McMillan, Equilibrium properties of the two-dimensional random (+ or-j) ising model, J. Phys. C **16**, 7027 (1983).
[22] Z. Zhu and H. G. Katzgraber, Do tensor renormalization group methods work for frustrated spin systems? arXiv:1903.07721.
[23] L. Vanderstraeten, B. Vanhecke, and F. Verstraete, Residual entropies for three-dimensional frustrated spin systems with tensor networks, Phys. Rev. E **98**, 042145 (2018).
[24] B. Vanhecke, J. Colbois, L. Vanderstraeten, F. Mila, and F. Verstraete, Relaxing frustration in classical spin systems, Phys. Rev. Research **3**, 013041 (2021).
[25] A. García-Sáez and J. I. Latorre, An exact tensor network for the 3SAT problem, Quantum Inf. Comput. **12**, 283292 (2012).
[26] J. D. Biamonte, J. Morton, and J. Turner, Tensor network contractions for #SAT, J. Stat. Phys. **160**, 1389 (2015).
[27] S. Kourtis, C. Chamon, E. R. Mucciolo, and A. E. Ruckenstein, Fast counting with tensor networks, SciPost Phys. **7**, 060 (2019).
[28] N. de Beaudrap, A. Kissinger, and K. Meichanetzidis, Tensor network rewriting strategies for satisfiability and counting, arXiv:2004.06455.
[29] D. Maclagan and B. Sturmfels, *Introduction to Tropical Geometry* (American Mathematical Society, Providence, 2015), Vol. 161.
[30] F. R. Kschischang, B. J. Frey, and H. A. Loeliger, Factor graphs and the sum-product algorithm, IEEE Trans. Inf. Theory **47**, 498 (2001).
[31] S. M. Aji and R. J. McEliece, The generalized distributive law, IEEE Trans. Inf. Theory **46**, 325 (2000).
[32] F. Obermeyer, E. Bingham, M. Jankowiak, D. Phan, and J. P. Chen, Functional tensors for probabilistic programming, arXiv:1910.10775.
[33] A. M. Rush, Torch-struct: Deep structured prediction library, arXiv:2002.00876.
[34] M. Mezard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, New York, 2009).
[35] N. G. Kingsbury and P. J. W. Rayner, Digital filtering using logarithmic arithmetic, Electron. Lett. **7**, 56 (1971).
[36] P. Zhang, Y. Zeng, and H. Zhou, Stability analysis on the finite-temperature replica-symmetric and first-step replica-symmetry-broken cavity solutions of the random vertex cover problem, Phys. Rev. E **80**, 021122 (2009).
[37] R. Marinescu and R. Dechter, Counting the optimal solutions in graphical models, Adv. Neural Inf. Process. Syst. **32**, 12091 (2019), https://proceedings.neurips.cc/paper/2019/file/fc2e6a440b94f64831840137698021e1-Paper.pdf.
[38] S. L. Lauritzen and D. J. Spiegelhalter, Local computations with probabilities on graphical structures and their application to expert systems, J. R. Stat. Soc. Ser. B **50**, 157 (1988).
[39] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Computational Complexity of Projected Entangled Pair States, Phys. Rev. Lett. **98**, 140506 (2007).
[40] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, T. Magerlein, E. Solomonik, E. W. Draeger, E. T. Holland, and R. Wisnieff, Breaking the 49-qubit barrier in the simulation of quantum circuits, arXiv:1710.05867.
[41] E. S. Fried, N. P. Sawaya, Y. Cao, I. D. Kivlichan, J. Romero, and A. Aspuru-Guzik, QTOrch: The quantum tensor contraction handler, PLoS One **13**, 1 (2018).
[42] E. F. Dumitrescu, A. L. Fisher, T. D. Goodrich, T. S. Humble, B. D. Sullivan, and A. L. Wright, Benchmarking tree-width as a practical component of tensor network simulations, PLoS One **13**, e0207827 (2018).
[43] J. M. Dudek, L. Dueñas-Osorio, and M. Y. Vardi, Efficient contraction of large tensor networks for weighted model counting through graph decompositions, arXiv:1908.04381.
[44] B. Villalonga, D. Lyakh, S. Boixo, H. Neven, T. S. Humble, R. Biswas, E. G. Rieffel, A. Ho, and S. Mandrà, Establishing the quantum supremacy frontier with a 281 Pflop/s simulation, Quantum Sci. Technol. **5**, 034003 (2020).
[45] F. Schindler and A. S. Jermyn, Algorithms for tensor network contraction ordering, arXiv:2001.08063.
[46] R. Schutski, D. Kolmakov, T. Khakhulin, and I. Oseledets, Simple heuristics for efficient parallel tensor contraction and quantum circuit simulation, Phys. Rev. A **102**, 062614 (2020).
[47] J. Gray and S. Kourtis, Hyper-optimized tensor network contraction, arXiv:2002.01935.
[48] C. Huang, F. Zhang, M. Newman, J. Cai, X. Gao, Z. Tian, J. Wu, H. Xu, H. Yu, B. Yuan, M. Szegedy, Y. Shi, and J. Chen, Classical simulation of quantum supremacy circuits, arXiv:2005.06787.
[49] In cases of the degenerated ground state, the approach gives one out of many ground state configurations. The particular configuration is selected by the default implementation of the maximum function, which returns the first argument when the two arguments are equal. One could obtain other degenerate solutions by changing this default behavior.
[50] H.-J. Liao, J.-G. Liu, L. Wang, and T. Xiang, Differentiable Programming Tensor Networks, Phys. Rev. X **9**, 031041 (2019).
[51] https://matbesancon.github.io/post/2020-01-23-discrete-diff/.
[52] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, Automatic differentiation in machine learning:

A survey, J. Mach. Learn. **18**, 1 (2018), https://jmlr.org/papers/v18/17-468.html.
[53] J. Revels, M. Lubin, and T. Papamarkou, Forward-mode automatic differentiation in Julia, arXiv:1607.07892.
[54] J.-G. Liu and T. Zhao, Differentiate everything with a reversible programming language, arXiv:2003.04617.
[55] X.-Z. Luo, J.-G. Liu, P. Zhang, and L. Wang, Yao.jl: Extensible, efficient framework for quantum algorithm design, Quantum **4**, 341 (2020).
[56] In general, it is always possible to map the tensor network contraction to a quantum circuit simulation by possibly introducing extra ancilla qubits.
[57] A. Selby, Efficient subgraph-based sampling of Ising-type models with frustration, arXiv:1409.3934.
[58] K. Jałowiecki, M. M. Rams, and B. Gardas, Brute-forcing spin-glass problems with CUDA, arXiv:1904.03621.
[59] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, Evidence for quantum annealing with more than one hundred qubits, Nat. Phys. **10**, 218 (2014).
[60] J. Vannimenus and G. Toulouse, Theory of the frustration effect. II. Ising spins on a square lattice, J. Phys. C **10**, L537 (1977).
[61] I. Morgenstern and K. Binder, Magnetic correlations in two-dimensional spin-glasses, Phys. Rev. B **22**, 288 (1980).
[62] J.-S. Wang and R. H. Swendsen, Low-temperature properties of the $\pm j$ Ising spin glass in two dimensions, Phys. Rev. B **38**, 4840 (1988).
[63] B. A. Berg and T. Celik, New Approach to Spin-Glass Simulations, Phys. Rev. Lett. **69**, 2292 (1992).
[64] L. Saul and M. Kardar, Exact integer algorithm for the two-dimensional $\pm j$ Ising spin glass, Phys. Rev. E **48**, R3221 (1993).
[65] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, and H. Neven, Simulation of low-depth quantum circuits as complex undirected graphical models, arXiv:1712.05384.
[66] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, Nat. Phys. **14**, 595 (2018).
[67] M. Mézard and G. Parisi, The Bethe lattice spin glass revisited, Eur. Phys. J. B **20**, 217 (2001).
[68] M. Mézard and G. Parisi, The cavity method at zero temperature, J. Stat. Phys. **111**, 1 (2003).
[69] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.126.090506 for the mapping of tensor networks to quantum circuits, reversible programming, and more experimental results on (Ising and Potts) spin glasses on square lattices, cubic lattice, and random graphs.
[70] C. De Simone, M. Diehl, M. Jünger, P. Mutzel, G. Reinelt, and G. Rinaldi, Exact ground states of Ising spin glasses: New experimental results with a branch-and-cut algorithm, J. Stat. Phys. **80**, 487 (1995).
[71] C. De Simone, M. Diehl, M. Jünger, P. Mutzel, G. Reinelt, and G. Rinaldi, Exact ground states of two-dimensional $\pm$ j Ising spin glasses, J. Stat. Phys. **84**, 1363 (1996).
[72] A. Percus, G. Istrate, and C. Moore, *Computational Complexity and Statistical Physics* (Oxford University Press, New York, 2006).
[73] B. Ghaddar, M. F. Anjos, and F. Liers, A branch-and-cut algorithm based on semidefinite programming for the minimum k-partition problem, Ann. Oper. Res. **188**, 155 (2011).
[74] M. F. Anjos, B. Ghaddar, L. Hupp, F. Liers, and A. Wiegele, Solving k-way graph partitioning problems to optimality: The impact of semidefinite relaxations and the bundle method, in *Facets of Combinatorial Optimization* (Springer, New York, 2013), pp. 355–386.
[75] https://github.com/TensorBFS/TropicalTensors.jl.
[76] https://github.com/YingboMa/MaBLAS.jl.