

Tensor-Network Markov-Chain Monte Carlo: Efficient Sampling of Three-Dimensional Spin Glasses

Tao Chen^{1,2}, Jing Liu^{3,4,*}, Youjin Deng^{1,2,5,†} and Pan Zhang^{4,6,2,7,‡}

¹*Hefei National Laboratory for Physical Sciences at the Microscale and Department of Modern Physics, University of Science and Technology of China, Hefei 230026, China*

²*Hefei National Laboratory, University of Science and Technology of China, Hefei 230088, China*

³*School of Physical Science and Technology, Beijing University of Posts and Telecommunications, Beijing 100876, China*

⁴*Institute of Theoretical Physics, Chinese Academy of Sciences, Beijing 100190, China*

⁵*College of Physics, Guizhou University, Guiyang 550025, China*

⁶*School of Fundamental Physics and Mathematical Sciences, Hangzhou Institute for Advanced Study, UCAS, Hangzhou 310024, China*

⁷*School of Physical Sciences, University of Chinese Academy of Sciences, Beijing 100049, China*



(Received 28 September 2025; revised 15 June 2026; accepted 12 August 2026; published 9 September 2026)

Sampling equilibrium configurations of three-dimensional (3D) spin glasses with quenched disorder remains a fundamental challenge in statistical physics. The rugged energy landscape, pronounced critical slowing down, and intrinsic ergodicity breaking render standard Monte Carlo methods severely inefficient, particularly for large systems at low temperatures. In this Letter, we introduce the tensor-network Markov-chain Monte Carlo (TNMCMC) approach to address the issue. By formulating large-scale collective updates via tensor-network contractions on 2D lattice slices, TNMCMC yields unbiased samples of the Boltzmann distribution while substantially reducing autocorrelation times relative to conventional MCMC approaches. Numerical benchmarks on 3D spin glasses demonstrate that the method systematically mitigates critical slowing down. Furthermore, application to three- and four-state 3D Potts models shows that TNMCMC can traverse the exponential free-energy barriers of first-order phase transitions that typically hinder local MC updates. These results establish TNMCMC as a promising numerical framework capable of addressing important 3D problems in statistical physics and beyond, such as the precise computation of low-temperature thermodynamic properties of three-dimensional spin glasses and the decoding of quantum error-correcting codes with a three-dimensional topology.

DOI: [10.1103/96lr-z7tb](https://doi.org/10.1103/96lr-z7tb)

Three-dimensional spin glasses combine lattice geometry, quenched disorder, and frustration, exhibiting a rugged energy landscape [1–3]. They are pivotal in statistical physics for providing profound insights into the nature of phase transitions [4], testing competing theories of replica symmetry breaking, benchmarking sampling and optimization algorithms, and providing the baseline for understanding glassy dynamics and aging phenomena. They remain one of the last unsolved challenges in statistical mechanics, and their equilibrium and dynamical properties remain analytically intractable despite half a century of effort.

For decades, Markov-chain Monte Carlo (MCMC) [5,6] (most commonly single-spin-flip Metropolis algorithms [7]) has been the standard workhorse for sampling and studying three-dimensional (3D) spin glasses. In this paradigm, one constructs a Markov chain whose stationary distribution is the Boltzmann distribution at a given temperature; successive configurations are generated by proposing local spin flips and accepting them according to the Metropolis-Hastings scheme. However, the standard single-spin-flip MCMC is severely inefficient at low temperatures due to the critical slowing down caused by the rugged energy landscape. The correlation time grows quickly with the system size and inverse temperature β , and the spin flips become locked in exponentially numerous metastable valleys, so the Markov chain needs exponentially many flips to reach equilibrium. Although advanced techniques such as parallel tempering and population annealing can accelerate the mixing of the Markov chain at low temperatures, it is still difficult to study large systems due to the rapid growth of the computational cost with system size.

An effective way to address the critical slowing down of MCMC is the *cluster update*, which updates a cluster

*Contact author: jing.liu@bupt.edu.cn

†Contact author: yjdeng@ustc.edu.cn

‡Contact author: panzhang@itp.ac.cn

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

(or block) of spins, creating a configuration that is much less correlated with the previous configuration than the case of single-spin flip. The examples include the celebrated Swendsen-Wang and Wolff cluster algorithms [8,9], which can dramatically enhance efficiency for ferromagnetic systems, due to their intrinsic connection to the Fortuin-Kasteleyn representation [10]. Although there have been recent efforts toward cluster-type algorithms for spin glasses [11–13], and a variant of Houdayer’s algorithm [14] has been reported to alleviate critical slowing down in 3D [15], it remains an open question whether a cluster method can correctly capture spin-spin correlations in spin-glass systems. A distinct strategy employs variational autoregressive networks [16,17] to generate global configurations via neural networks, though these require costly training by minimizing the variational free energy.

Here, we introduce the tensor-network Markov-chain Monte Carlo (TNMCMC) to tackle the sampling challenge in 3D spin glasses. By mapping 2D slices of the 3D lattice into 2D tensor networks, we generate whole-slice spin configurations via tensor-network contraction, delivering large-scale collective updates that leap over local minima and suppress critical slowing down. The bias of generated samples induced by 2D tensor-network contractions is eliminated through a Metropolis-Hastings algorithm, ensuring detailed balance. To benchmark the efficiency of our approach, we apply it to the 3D Edwards-Anderson (EA) spin glass. The results show that TNMCMC dramatically suppresses critical slowing down at low temperatures, reaching equilibrium many orders of magnitude faster than standard MCMC in both the number of MC steps and computational time. The advantage of our approach increases quickly as the system size increases and the temperature decreases. We further apply TNMCMC to the 3D three- and four-state Potts models, which exhibit first-order phase transitions that are challenging for standard MCMC, and we show that TNMCMC can overcome the exponential tunneling barrier efficiently. In the following text, we will first introduce the TNMCMC approach, then present the numerical results in 3D spin glasses and the Potts model in detail, and discuss the further generalizations of TNMCMC.

TNMCMC—Consider the Boltzmann distribution of n spins,

$$P(\mathbf{s}) = \frac{1}{Z} e^{-\beta E(\mathbf{s})}, \quad (1)$$

where $\mathbf{s} \in \{+1, -1\}^n$ is a configuration, β is the inverse temperature, $E(\mathbf{s})$ is the energy function, and $Z = \sum_{\mathbf{s}} e^{-\beta E(\mathbf{s})}$ denotes the partition function. Sampling from the Boltzmann distribution belongs to the computational class of #P-hard problems [18], and no exact algorithms are known for solving the general problem in polynomial time.

In standard MCMC, a Markov chain of configurations is constructed for generating samples. At each step, given the current state $\mathbf{s}$, a candidate $\mathbf{s}'$ is drawn from a proposal distribution $q(\mathbf{s}')$ and accepted with probability,

$$A(\mathbf{s} \rightarrow \mathbf{s}') = \min \left\{ 1, \frac{q(\mathbf{s}) \times e^{-\beta E(\mathbf{s}')}}{q(\mathbf{s}') \times e^{-\beta E(\mathbf{s})}} \right\}. \quad (2)$$

The Metropolis algorithm [7] uses a candidate $\mathbf{s}'$ that differs from $\mathbf{s}$ by only one spin. This method has a rather high acceptance probability but cannot cross a large energy barrier, and thus it may need exponentially many steps to decorrelate in the presence of free-energy barriers.

To address this issue, we propose to update entire 2D slices of the 3D lattice using tensor networks. At each step, spins are partitioned into sampling spins $\mathbf{s}$ on selected slices and fixed spins σ elsewhere, as illustrated in Fig. 1 for two partitioning schemes. The conditional distribution $P(\mathbf{s}|\sigma)$ is approximately factorized into a product of autoregressive conditionals and computed via tensor-network contraction [19–21], yielding a proposal distribution from which




FIG. 1. Schematic of the tensor-network Markov-chain Monte Carlo (TNMCMC) for efficient sampling of 3D spin glasses. At each step of sampling, spins are separated into two groups: sampling spins with configuration $\mathbf{s}$ (in the blue-shaded slices of the 3D lattice) and fixed spins with configuration σ (red spins). We present two sampling-spin schemes: (a) the spins to be sampled lie on alternating 2D slices across the 3D lattice, separated by the fixed spins; (b) the sampling spins form pairs of two adjacent slices that can be merged into a single “super-spin” layer. For each scheme, the conditional probability $P(\mathbf{s}|\sigma)$ is computed approximately with a tensor network, and configurations $\mathbf{s}$ are drawn and then resampled via the Metropolis-Hastings rule [Eq. (2)] to remove the bias introduced by the approximate tensor-network contractions.

candidate configurations $\mathbf{s}'$ are drawn and accepted via Eq. (2). This Metropolis-Hastings step rigorously removes any bias introduced by the approximate contractions. The 2D-slice geometry is chosen to optimize the trade-off between cluster size and contraction accuracy. Single-spin flips offer no collective moves, and contracting the full 3D lattice is prohibitive. By contrast, sampling 2D slices exploits efficient tensor-network algorithms [22,23] to execute large-scale, collective spin reconfigurations. Through sequential updates across alternating slice orientations, the method acts as an anisotropic, nonlocal accelerator. While incapable of reconstructing entire droplets in a single step, it reshapes 3D domain interfaces via successive accepted rearrangements. An outline of the sampling procedure using conditional probabilities is provided in the End Matter, and the detailed algorithms for the two 2D-slice schemes are given in Supplemental Material [24].

Sampling 3D spin glasses—To test our approach, we conduct numerical experiments for a prototypical spin-glass model, the 3D EA model [1], on an $L \times L \times L$ cubic lattice with free boundary conditions, with energy function

$$E(\mathbf{s}) = -\sum_{\langle ij \rangle} J_{ij} s_i s_j, \quad (3)$$

where $s_i \in \{-1, +1\}$, and the coupling J_{ij} takes 1 with probability p and -1 with probability $1 - p$. For $p = 0.5$, the system exhibits a spin-glass transition at inverse temperature $\beta_c = 1/T_c = 0.9075(24)$ [27]. Above β_c , the system enters a glassy phase characterized by a rugged energy landscape and slow dynamics, and standard MCMC methods are hampered by relaxation times that diverge exponentially with system size and inverse temperature.

To assess the sampling efficiency, we compute the integrated autocorrelation time $\tau_{\text{int},|q|}$ of the spin-glass overlap $q = (1/L^3) \sum_i s_i^{(1)} s_i^{(2)}$, where $s_i^{(1)}$ and $s_i^{(2)}$ are spin configurations from two independent replicas. Figure 2(a) shows $\tau_{\text{int},|q|}$ and the computational cost per independent sample $t_{\text{int},|q|} = \tau_{\text{int},|q|} \times t_s$ for $L = 8$, where TNMCMC achieves orders-of-magnitude savings over single-spin Metropolis in both metrics. For larger systems, evaluating $\tau_{\text{int},|q|}$ becomes computationally prohibitive.




FIG. 2. Results for the 3D Edwards-Anderson spin glasses. (a) Integrated autocorrelation time $\tau_{\text{int},|q|}$ and computational cost per independent sample $t_{\text{int},|q|} = \tau_{\text{int},|q|} \times t_s$ for $L = 8$, with t_s the wall-clock time per sweep. The vertical dashed line marks the spin-glass transition $\beta_c = 1/T_c = 0.9075(24)$. (b) First-passage time to zero magnetization $\tau_{\text{fp},M}$ and corresponding wall-clock time $t_{\text{fp},M}$ as a function of L at $\beta = 1.5$. The solid line is a fit to $\tau_0 \exp(cL^\psi)$; see End Matter for details. (c) $\tau_{\text{int},|q|}$ as a function of β for $L = 12$, comparing four protocols: single-spin Metropolis, TNMCMC-2, Metropolis + PT, and TNMCMC-2 + PT, with 21 replicas in PT simulations. (d)–(f) Normalized autocorrelation function $C_q^2(t)$ of the squared overlap for $L = 16$ at $\beta \approx 0.769, 0.910$, and 1.200 , comparing Metropolis (blue) and TNMCMC (red).

We, therefore, introduce the first-passage time to zero magnetization $\tau_{\text{fp},M}$, starting from an all-up configuration, which measures the characteristic escape time from metastable states (see Supplemental Material [24] for rigorous definitions). Figure 2(b) presents $\tau_{\text{fp},M}$ as a function of L at $\beta = 1.5$. TNMCMC dramatically outperforms Metropolis, with the gap widening exponentially as L increases. The scaling form $\tau_{\text{fp},M} \approx \tau_0 \exp(cL^\psi)$ and the fitted barrier exponents are reported in the End Matter.

Since parallel tempering (PT) [28] combined with single-spin Metropolis is the standard baseline for spin-glass simulations, Fig. 2(c) compares four protocols for $L = 12$ on equal footing: single-spin Metropolis, standalone TNMCMC-2, Metropolis + PT, and TNMCMC-2+PT, with all PT simulations employing 21 replicas (see End Matter for a description of the TNMCMC + PT combination). TNMCMC + PT yields the lowest $\tau_{\text{int},|q|}$ across the entire temperature range, demonstrating that the tensor-network proposal complements replica exchange. Figures 2(d)–2(f) display the normalized autocorrelation function $C_{q^2}(t) = \{[\langle q^2(t_0 + t)q^2(t_0) \rangle - \langle q^2 \rangle^2] / [\langle (q^2)^2 \rangle - \langle q^2 \rangle^2]\}$ of the squared overlap for $L = 16$ at $\beta \simeq 0.769, 0.910$, and 1.200 . All correlation data are measured under well-equilibrated equilibrium states after sufficient thermalization. Here, Metropolis exhibits a slow, temperature-dependent relaxation tail, whereas TNMCMC rapidly decorrelates within a short time window, consistent with its much smaller integrated autocorrelation time. Detailed overlap-based equilibration diagnostics are provided in Supplemental Material [24].

Sampling the 3D Potts model—The 2D-slice update via tensor networks is an approach generally applicable to 3D systems. Here, we give an example using the

three-state Potts model in 3D, with the energy function defined as

$$E(\sigma) = -J \sum_{\langle i,j \rangle} \delta(\sigma_i, \sigma_j), \quad (4)$$

where $\sigma_i \in \{1, 2, 3\}$, $J > 0$ is the ferromagnetic coupling strength, and δ is the Kronecker delta. This model is known to exhibit a weak first-order phase transition [29], where conventional MCMC algorithms typically struggle [29,30]. We conduct numerical experiments on the Potts model up to system size $64 \times 64 \times 64$ and present the scaling of the first-passage time to zero magnetization $\tau_{\text{fp},M}$ at its critical inverse temperature $\beta_c = 0.550565$ [29]. In this case, we employ periodic boundary conditions and use TNMCMC-1 only. At each step, a random slice is selected from the $x-y$, $y-z$, or $z-x$ planes, and the spins on a randomly chosen line are held fixed.

As shown in Fig. 3(a), local Metropolis updates yield an exponential divergence of $\tau_{\text{fp},M}$ with system size L , reflecting the algorithm's inability to traverse the interfacial free-energy barrier. Conversely, TNMCMC-1 dramatically suppresses this scaling, achieving substantially smaller first-passage times. The energy time series [Fig. 3(b)] further reveals that the Metropolis dynamics remains metastably trapped in the ordered phase, sampling only local domain fluctuations. In stark contrast, TNMCMC-1 exhibits rapid tunneling between the ordered and disordered phases, successfully sampling mixed-phase intermediate states as evidenced by the coexistence histogram [Fig. 3(c)]. These results confirm that the nonlocal, tensor-network-guided updates efficiently traverse rugged free-energy landscapes. Analogous success in the four-state Potts model, which features a stronger first-order transition




FIG. 3. Results for the 3D three-state Potts model with first-order transition. (a) The first-passage time to zero magnetization, $\tau_{\text{fp},M}$, obtained by the conventional Metropolis algorithm and TNMCMC-1 (with bond dimension $\chi = 3$), at the critical $\beta_c = 0.550565$. Each point is the statistical average, with the number of instances ranging from 15 to 1024. (b) Energy of samples obtained using TNMCMC-1 (red) and Metropolis (blue) for a lattice with $L = 48$ at β_c , where the system exhibits two energy barriers as illustrated by the histogram of TNMCMC-1 configurations in (c). Notice that, while the conventional Metropolis becomes trapped into one energy valley, TNMCMC jumps back and forth between two energy barriers, efficiently traversing the exponential barriers of the first-order transition. The bond dimension employed in TNMCMC-1 is 3.

and a pronounced double-peaked energy distribution, is detailed in Supplemental Material [24].

Conclusions and discussion—We have introduced TNMCMC, a general-purpose algorithm for simulating challenging 3D statistical physics systems efficiently. TNMCMC integrates tensor-network algorithms with MCMC sampling. By slicing the 3D lattice into 2D layers and sampling spins in each layer with tensor networks, the method delivers nonlocal, collective moves that leap across rugged energy landscapes. A final Metropolis-Hastings acceptance step rigorously enforces detailed balance, ensuring both ergodicity and correctness, while sidestepping the infeasible cost and large error of fully contracting 3D tensor networks. We have demonstrated its efficiency on the 3D Edwards-Anderson spin glass and the three- and four-state 3D Potts models, showing that TNMCMC mitigates critical slowing down in the glassy phase and overcomes the free-energy barriers of a first-order phase transition.

TNMCMC is designed to complement highly optimized conventional MCMC routines. The nonlocal updates can be seamlessly combined with standard single-spin sweeps or serve as the core update kernel within parallel tempering (PT). Crucially, our numerical experiments confirm that embedding TNMCMC within PT systematically accelerates low-temperature equilibration relative to standard Metropolis-PT baselines. Because the 2D tensor-network sampling is model agnostic, this methodology generalizes to a broad class of frustrated 3D classical models and 2D quantum many-body systems mapped to 3D classical representations [31–33].

Looking ahead, the efficiency and generality of TNMCMC open the door to a range of important problems that remain difficult for conventional numerical approaches. Most directly, with increased computational resources such as large-scale graphics processing unit (GPU) clusters, we expect the method to enable precise computation of low-temperature thermodynamic properties of three-dimensional spin glasses. Moreover, the decoding of quantum codes, such as surface codes, is an important problem in quantum computing. With circuit-level noise, the topology of the underlying detector error model becomes 3D, and the decoding problem is closely related to the statistical mechanics of spin glasses; our approach, thus, offers a potentially new and precise way of tackling this problem.

While the algorithm already performs strongly in 3D systems, the underlying two-dimensional tensor-network contractions and sampling implementation remain suboptimal. We anticipate significant speedups through further engineering and GPU acceleration [34–37]. In the current version, tensor-network contractions are performed frequently since the to-be-updated slices are subject to effective external fields from their adjacent layers of fixed spins. One might generate (and dynamically update) a long list of “typical” external fields so that tensor-network contractions are only needed occasionally. This would introduce further

approximations and reduce acceptance probabilities, requiring a trade-off between accuracy and efficiency.

Acknowledgments—This work is supported by Projects No. 12325501, No. 12247104, No. 12275263, and No. 12405047 of the National Natural Science Foundation of China, the Innovation Program for Quantum Science and Technology (under Grant No. 2021ZD0301900), and the Strategic Priority Research Program of Chinese Academy of Sciences, Grant No. XDB1680000. Y. D. is thankful for the support from the Natural Science Foundation of Fujian Province of China (under Grant No. 2023J02032).

Data availability—The numerical data supporting the findings of this article and the source code used to generate the simulations are openly available from Ref. [38].

-
- [1] S. F. Edwards and P. W. Anderson, Theory of spin glasses, *J. Phys. F* **5**, 965 (1975).
 - [2] K. Binder and A. P. Young, Spin glasses: Experimental facts, theoretical concepts, and open questions, *Rev. Mod. Phys.* **58**, 801 (1986).
 - [3] A. Young, Numerical simulations of spin glasses: Methods and some recent results, in *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology*, edited by M. Ferrario, G. Ciccotti, and K. Binder (Springer, Berlin, Heidelberg, 2006), Vol. 2, pp. 31–44.
 - [4] H. Nishimori and G. Ortiz, *Elements of Phase Transitions and Critical Phenomena* (Oxford University Press, New York, 2010).
 - [5] M. Newman and G. Barkema, *Monte Carlo Methods in Statistical Physics* (Clarendon Press, Oxford, 1999).
 - [6] D. P. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics*, 4th ed. (Cambridge University Press, Cambridge, England, 2014).
 - [7] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.* **21**, 1087 (1953).
 - [8] R. H. Swendsen and J.-S. Wang, Nonuniversal critical dynamics in Monte Carlo simulations, *Phys. Rev. Lett.* **58**, 86 (1987).
 - [9] U. Wolff, Collective Monte Carlo updating for spin systems, *Phys. Rev. Lett.* **62**, 361 (1989).
 - [10] C. M. Fortuin and P. W. Kasteleyn, On the random-cluster model: I. Introduction and relation to other models, *Physica (Amsterdam)* **57**, 536 (1972).
 - [11] G. A. Miranda, M. Zheng, P. Charbonneau, A. Coniglio, L. F. Cugliandolo, and M. Tarzia, Percolation and criticality of systems with competing interactions on Bethe lattices: Limitations and potential strengths of cluster schemes, *Phys. Rev. E* **114**, 014114 (2026).
 - [12] L. Münster and M. Weigel, Cluster percolation in the two-dimensional Ising spin glass, *Phys. Rev. E* **107**, 054103 (2023).
 - [13] L. Münster and M. Weigel, Cluster percolation in the three-dimensional $\pm J$ random-bond Ising model, *Phys. Rev. E* **113**, 024139 (2026).

- [14] J. Houdayer, A cluster Monte Carlo algorithm for 2-dimensional spin glasses, *Eur. Phys. J. B* **22**, 479 (2001).
- [15] Z. Zhu, A. J. Ochoa, and H. G. Katzgraber, Efficient cluster algorithm for spin glasses in any space dimension, *Phys. Rev. Lett.* **115**, 077201 (2015).
- [16] D. Wu, L. Wang, and P. Zhang, Solving statistical mechanics using variational autoregressive networks, *Phys. Rev. Lett.* **122**, 080602 (2019).
- [17] L. M. Del Bono, F. Ricci-Tersenghi, and F. Zamponi, Performance of machine-learning-assisted Monte Carlo in sampling from simple statistical physics models, *Phys. Rev. E* **112**, 045307 (2025).
- [18] C. Moore and S. Mertens, *The Nature of Computation* (Oxford University Press, New York, 2011).
- [19] M. Frías Pérez, M. Mariën, D. Pérez García, M. C. Bañuls, and S. Iblisdir, Collective Monte Carlo updates through tensor network renormalization, *SciPost Phys.* **14**, 123 (2023).
- [20] T. Chen, E. Guo, W. Zhang, P. Zhang, and Y. Deng, Tensor network Monte Carlo simulations for the two-dimensional random-bond Ising model, *Phys. Rev. B* **111**, 094201 (2025).
- [21] 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).
- [22] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Ann. Phys. (Amsterdam)* **349**, 117 (2014).
- [23] S.-J. Ran, E. Tirrito, C. Peng, X. Chen, L. Tagliacozzo, G. Su, and M. Lewenstein, *Tensor Network Contractions: Methods and Applications to Quantum Many-Body Systems* (Springer International Publishing, New York, 2020).
- [24] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/96lr-z7tb> for details of the TNMCMC algorithms, performance and equilibration diagnostics, and additional results for three-dimensional statistical models, which includes Refs. [25,26].
- [25] M. Levin and C. P. Nave, Tensor renormalization group approach to two-dimensional classical lattice models, *Phys. Rev. Lett.* **99**, 120601 (2007).
- [26] C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, New York, 2006).
- [27] M. Baity-Jesi *et al.* (Janus Collaboration), Critical parameters of the three-dimensional Ising spin glass, *Phys. Rev. B* **88**, 224416 (2013).
- [28] K. Hukushima and K. Nemoto, Exchange Monte Carlo method and application to spin glass simulations, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).
- [29] W. Janke and R. Villanova, Three-dimensional 3-state Potts model revisited with new techniques, *Nucl. Phys. B* **489**, 679 (1997).
- [30] B. A. Berg and T. Neuhaus, Multicanonical algorithms for first order phase transitions, *Phys. Lett. B* **267**, 249 (1991).
- [31] W.-Y. Liu, Y.-Z. Huang, S.-S. Gong, and Z.-C. Gu, Accurate simulation for finite projected entangled pair states in two dimensions, *Phys. Rev. B* **103**, 235155 (2021).
- [32] W.-Y. Liu, H. Zhai, R. Peng, Z.-C. Gu, and G. K.-L. Chan, Accurate simulation of the Hubbard model with finite fermionic projected entangled pair states, *Phys. Rev. Lett.* **134**, 256502 (2025).
- [33] D. Layden, G. Mazzola, R. V. Mishmash, M. Motta, P. Wocjan, J.-S. Kim, and S. Sheldon, Quantum-enhanced Markov chain Monte Carlo, *Nature (London)* **619**, 282 (2023).
- [34] M. Weigel, Performance potential for simulating spin models on GPU, *J. Comput. Phys.* **231**, 3064 (2012).
- [35] M. Bisson, M. Bernaschi, M. Fatica, N. G. Fytas, I. González-Adalid Pemartín, V. Martín-Mayor, and A. Vasilopoulos, Massive-scale simulations of 2D Ising and Blume-Capel models on rack-scale multi-GPU systems, *Comput. Phys. Commun.* **315**, 109690 (2025).
- [36] M. Bisson, A. Vasilopoulos, M. Bernaschi, M. Fatica, N. G. Fytas, I. G.-A. Pemartín, and V. Martín-Mayor, Universal exotic dynamics in critical mesoscopic systems: Simulating the square root of Avogadro's number of spins, *Phys. Rev. Res.* **7**, 033218 (2025).
- [37] T. Chen, J. Zhang, J. Liu, Y. Deng, and P. Zhang, Batchtnmc: Efficient sampling of two-dimensional spin glasses using tensor network Monte Carlo, *arXiv:2509.19006*.
- [38] T. Chen, J. Liu, Y. Deng, and P. Zhang, TNMCMC: Tensor-network Markov-chain Monte Carlo, GitHub, <https://github.com/Fermichen99/TNMCMC> (2026).
- [39] J.-G. Liu, L. Wang, and P. Zhang, Tropical tensor network for ground states of spin glasses, *Phys. Rev. Lett.* **126**, 090506 (2021).
- [40] K. Okunishi, T. Nishino, and H. Ueda, Developments in the tensor network—From statistical mechanics to quantum entanglement, *J. Phys. Soc. Jpn.* **91**, 062001 (2022).
- [41] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [42] T. Nishino and K. Okunishi, A density matrix algorithm for 3D classical models, *J. Phys. Soc. Jpn.* **67**, 3066 (1998).
- [43] Q. N. Chen, M. P. Qin, J. Chen, Z. C. Wei, H. H. Zhao, B. Normand, and T. Xiang, Partial order and finite-temperature phase transitions in Potts models on irregular lattices, *Phys. Rev. Lett.* **107**, 165701 (2011).
- [44] Z. Y. Xie, H. C. Jiang, Q. N. Chen, Z. Y. Weng, and T. Xiang, Second renormalization of tensor-network states, *Phys. Rev. Lett.* **103**, 160601 (2009).
- [45] X. F. Liu, Y. F. Fu, W. Q. Yu, J. F. Yu, and Z. Y. Xie, Variational corner transfer matrix renormalization group method for classical statistical models, *Chin. Phys. Lett.* **39**, 067502 (2022).
- [46] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete, Matrix product states and projected entangled pair states: Concepts, symmetries, theorems, *Rev. Mod. Phys.* **93**, 045003 (2021).
- [47] T. Xiang, *Density Matrix and Tensor Network Renormalization* (Cambridge University Press, Cambridge, England, 2023).
- [48] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (Amsterdam)* **326**, 96 (2011).

End Matter

Sampling using conditional probabilities and tensor-network contraction—In the TNMCMC approach, the joint distribution of sampling spins $\mathbf{s}$ conditioned on the fixed spins σ is factorized into a product of conditional probabilities,

$$P(\mathbf{s}|\sigma) = \prod_i P(s_i|\mathbf{s}_{<i}, \sigma), \quad (\text{A1})$$

where $\mathbf{s}_{<i}$ denotes spins already sampled and $\mathbf{s}_{>i}$ those remaining. Each conditional probability is given by

$$P(s_i|\mathbf{s}_{<i}, \sigma) = \frac{Z(s_i, \mathbf{s}_{<i}, \sigma)}{\sum_{s_i} Z(s_i, \mathbf{s}_{<i}, \sigma)}, \quad (\text{A2})$$

with $Z(s_i, \mathbf{s}_{<i}, \sigma) = \sum_{\mathbf{s}_{>i}} e^{-\beta E(s_i, \mathbf{s}_{<i}, \mathbf{s}_{>i}, \sigma)}$ computed via tensor-network contraction [21,39]. When computed exactly, the acceptance probability [Eq. (2)] is unity. However, exact contraction is feasible only for small systems, as the intermediate bond dimension grows exponentially with system size. For practical system sizes, approximate contractions with a fixed bond dimension are employed, introducing errors that reduce the acceptance probability [20,22,23,40–48]. The detailed algorithms for the two 2D-slice schemes (TNMCMC-1 and TNMCMC-2) and the choice of bond dimension are presented in Supplemental Material [24].

Scaling form and barrier exponents—At an inverse temperature $\beta = 1.5$, well above the critical inverse temperature $\beta_c \approx 0.9$, the first-passage time to zero magnetization $\tau_{\text{fp},M}$ grows exponentially with system size L , consistent with activated dynamics in glassy systems. The data are well described by the scaling form,

$$\tau_{\text{fp},M} \approx \tau_0 \exp(cL^\psi), \quad (\text{B1})$$

where ψ is the barrier exponent that characterizes the nature of the free-energy barriers. A fit of the numerical data to this form yields $\psi \approx 0.56$ for the Metropolis algorithm, $\psi \approx 0.16$ for TNMCMC-1, and $\psi \approx 0.0051$ for TNMCMC-2, indicating that TNMCMC substantially

reduces the effective barrier exponent. The systematically lower autocorrelation times for TNMCMC-2 arise because it updates spins in composite two-layer units, enabling larger configuration changes per Monte Carlo step compared to the single-layer updates of TNMCMC-1. We have also performed experiments at other temperatures and observed qualitatively similar results. For example, at $\beta = 1.0$, an inverse temperature slightly above $\beta_c \approx 0.9$, TNMCMC can successfully reach the zero-magnetization state, enabling measurement of the first-passage time with system size up to $L = 64$ with $\tau_{\text{fp},M}(\text{TNMCMC-2}) \approx 3.99 \times 10^3$, while $\tau_{\text{fp},M}(\text{Metropolis}) \approx 2.04 \times 10^6$.

Combining parallel tempering with TNMCMC—Parallel tempering (PT) [28] simulates R replicas of the system at a ladder of inverse temperatures $\beta_1 < \beta_2 < \dots < \beta_R$, with each replica evolving independently via MCMC at its assigned temperature. After a fixed number of MCMC sweeps, a swap of configurations between two adjacent replicas at β_k and β_{k+1} is proposed and accepted with probability

$$A_{\text{swap}} = \min\{1, \exp[(\beta_k - \beta_{k+1})(E_k - E_{k+1})]\}, \quad (\text{C1})$$

where $E_k = E(\mathbf{s}_k)$ denotes the energy of the configuration at β_k . This exchange mechanism allows low-temperature replicas to escape deep free-energy minima by borrowing configurations that equilibrated at higher temperatures, while the detailed balance condition is preserved. Since TNMCMC operates at a fixed temperature, it can serve as a drop-in replacement for the single-spin Metropolis kernel within the PT framework: each replica evolves via TNMCMC updates at its assigned temperature, while replica exchanges proceed according to the standard PT protocol. This combination leverages TNMCMC's ability to suppress critical slowing down within each replica, while PT overcomes free-energy barriers by exchanging configurations across temperatures, yielding a straightforward integration that preserves the standard exchange protocol without modification.