Topological Quantum Orders at Finite Temperature

Finite-Temperature Quantum Topological Order in Three Dimensions
Authors: Shu-Tong Zhou, Meng Cheng, Tibor Rakovszky, Curt von Keyserlingk, and Tyler D. Ellison
Phys. Rev. Lett. 135, 040402 (2025)

Recommended with a Commentary by Ashvin Vishwanath,
Harvard University

Background: A recurring conceptual question in quantum many-body physics is whether matter at nonzero temperature is genuinely quantum, or whether thermal fluctuations inevitably wash quantum structure into something that is, at least effectively, classical.

Topological order helps sharpen this question. At zero temperature, topologically ordered phases are paradigmatic examples of nontrivial phases with *long range entanglement*. Examples include fractional quantum Hall states and lattice models such as the toric code. The defining features of these gapped quantum phases include ground-state degeneracy on manifolds of nontrivial topology, and anyonic excitations. These features cannot be mimicked by any short-range-entangled (SRE) state - perhaps the most extreme example of a SRE state is just a product state, where eg. qubits on a lattice are put into individual states giving rise to a product state $|\psi\rangle = \prod_i |\phi_i\rangle$. More general SRE states are 'close' to such product states.

But what about thermal states? Here one typically deals not with a pure state but with a Gibbs density matrix $\rho_0 = e^{-\beta H}$. Are there phases in which the finite-temperature Gibbs state is unavoidably long-range entangled (LRE)? This immediately raises a technical question: how should we define long-range entanglement for mixed states, so as to meaningfully distinguish "nontrivial" density matrices from those that are merely short-range entangled (SRE)? Setting that question aside for the moment, and assuming we have a criterion that separates LRE from SRE mixed states, the central question becomes: do there exist thermal equilibrium states that are asymptotically distinct from all SRE states, in direct analogy with zero-temperature topologically ordered phases? This issue is not only conceptually important; it is a necessary condition for realizing a truly self-correcting quantum memory, a "quantum hard drive" that remains stable without active error correction.

Previous work painted a rather pessimistic picture for finite-temperature quantum topological order in physically realizable spatial dimensions ($d \leq 3$). While four-dimensional models such as the 4D toric code *can* support finite-temperature quantum memory and long-range entanglement in their thermal states, a combination of explicit constructions and no-go theorems has suggested that in 2D, and for broad classes of 3D models, thermal topological order is either absent or effectively classical [1]. In particular, the 3D toric code was shown to have nonzero-temperature thermal states that are SRE, even though the model exhibits a finite-temperature phase transition. The transition separates a low temperature phase with only few magnetic loop excitations from a high temperature phase where they proliferate. This transition separates classical phases [2], rather than distinct quantum-topologically ordered phases. *

The featured paper breaks this impasse. It answers the question "Can there be genuine quantum topological order at finite temperature in three dimensions?" in the affirmative. In does so by providing a concrete setting where finite-temperature equilibrium is provably and robustly quantum in its global entanglement structure. The paper focuses on a three-dimensional toric code, but with a literal twist. In 3D, the toric code comes in two variants - the traditional (bosonic) and twisted (fermionic) toric codes - with the latter being the model of interest here. It is defined on a cubic lattice with two-level systems (qubits) residing on the edges, governed by the Hamiltonian:

$$H = -\sum_v A_v - \sum_p B_p$$

where the terms are associated with vertices (v) and plaquettes (p) of the lattice. The vertex term A_v is simply the product of the Pauli operators σ^z on the six edges adjacent to a vertex. The plaquette terms are more complex for the fermionic toric code, (for the bosonic toric code this is nothing but the product of the σ^x Pauli operators around the plaquette) but again feature products of σ^x and σ^z on specific edges in the vicinity of a plaquette. Importantly, all vertex and plaquette operators in the Hamiltonian commute with one another, allowing for an exact solution. This model's particle like 'charge' excitations are fermions, unlike the bosons in the traditional 3D toric code, the model which was previously shown to be SRE at any finite temperature. The reader is left to extract the details from the recommended paper, but here are a few key physics points that should help demystify the proof.

The proof involves three steps. **(i)** The definition of SRE density matrices introduced in [1] is quite natural, as an incoherent sum of SRE states, i.e. construct a density matrix with states $|SRE_i\rangle$ with probability p_i. **(ii)** The reference density matrix $\rho_\emptyset$ is constructed as follows - start with the thermal state of the fermionic toric code model which may be viewed as a thermal soup of point 'charges' as well as closed loops. Now consider keeping the first sector, the thermal gas of point charges, but completely eliminating the latter, namely shrinking all loops to zero. This is the reference state $\rho_\emptyset$. A very simple argument suffices to show that this state *cannot* be approximated by any ρ_{SRE}. One exploits the fact that there are fermionic charges in the system and applies a series of unitary operations that isolates

*As a brief introduction to toric codes across dimensions: these models realize $\mathbb{Z}_2$ topological order in 2D and 3D. They support a $\mathbb{Z}_2$-valued charge excitation e and a corresponding $\mathbb{Z}_2$ flux excitation m, with nontrivial mutual statistics: transporting e around m produces a π phase. In 2D, both e and m are point-like anyons, and one can form their bound state $\varepsilon = e \times m$, which has fermionic statistics. In this sense, 2D $\mathbb{Z}_2$ topological order does not sharply distinguish "bosonic" versus "fermionic" charge. In 3D, by contrast, the flux becomes loop-like while the charge remains point-like, and the notion of a point-like bound state $e \times m$ is no longer available. As a result, the statistics of the point charges themselves becomes a genuine invariant: the bosonic and fermionic 3D toric codes correspond to distinct phases.

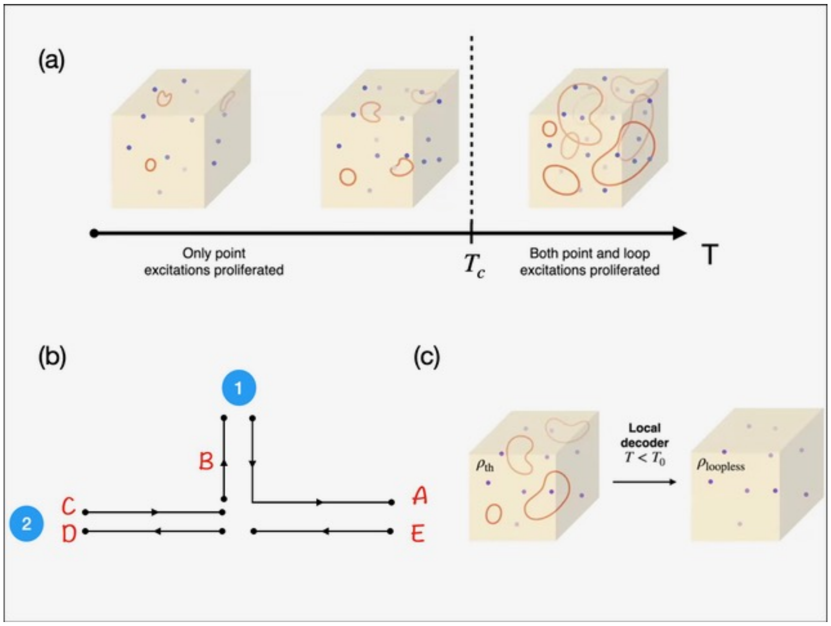

Figure 1: (a) Phase diagram of 3D toric code as a function of temperature. At T_c, loops of magnetic flux proliferate. (b) The T-junction setup to extract exchange statistics, the relative phase betweeen the sequence ACBED and AEBCD arises purely from exchange and (c) the cleanup step that effectively removes magnetic loops. Figures (a)-(c) are taken from Prof. Meng Cheng's seminar.

the fermionic exchange statistics. The unitary "string" operators that move the particles are shown in Figure 1(b) (adapted from [3]). By applying the operators in Sequence-1, ACEBD (i.e. A followed by C ... and ending with D) can be shown to interchange the two particles 1,2. On the other hand, the exact same set of steps in the Sequence-2, AEBCD, does no interchange. Because the two processes are otherwise identical, and the particles are assumed to remain well separated, all dynamical phases cancel. The only surviving distinction is a relative minus sign of exchanging a pair of fermions. Such a relative phase for a sequence of local string operators is *incompatible* with ρ_{SRE}. Thus, crudely speaking $\rho_\emptyset \neq \rho_{SRE}$. This intuition can be made precise in terms of the distance between the two kinds of density matrices [4]. Finally, **(iii)** we need to decide if the thermal density matrix of the fermionic toric code ρ_0 is closer to ρ_{SRE} or to $\rho_\emptyset$. This depends on the temperature which can be seen by a 'cleanup' procedure or decoder, a unitary operation that cleans up loops in the thermal state and makes it approach the $\rho_\emptyset$. Such a procedure should succeed as long as we are below the finite temperature transition T_c in Figure 1(a). For technical reasons the current proof is able to establish it upto a slightly lower temperature ($T_{\text{cleanup}} \sim 0.95T_c$) that depends on the details of the cleanup procedure. Thus on proving that ρ_0 is close to $\rho_\emptyset$ and hence far from ρ_{SRE} one has established the existence of a thermal state that has long range entanglement.[†]

[†]The slight discrepancy between the two temperatures T_{cleanup} and T_c is believed to be an artifact of the proof. An optimal decoder, which does the loop cleanup in the most efficient way, should place the boundary between quantum and classical states precisely at the phase transition.

Given the existence of a thermal quantum state one may hope this would be a candidate for a thermal quantum memory which retains coherent quantum information at finite temperature. Alas, this is not the case - like its bosonic counterpart the fermionic 3D toric code can only encode a classical bit below the transition temperature, exemplifying that finite temperature LRE is necessary but not sufficient for a thermal quantum memory.

Future directions and caveats: What I found appealing about this work is that it highlights an important and surprising phenomenon, 3D finite temperature topological order, hiding in plain sight within an extensively studied model-the fermionic toric code.

We close with a few remarks. **(a)** The proof hinges on a special symmetry of the model, generated by loop operators (so called 2-form symmetry). Indeed, the result can be viewed as a consequence of the anomalous (ultimately fermionic-statistics–related) character of this symmetry. Generic models in the same phase do not possess such a microscopic symmetry. An important open question is whether the result also holds if this symmetry emerges just in the low-energy limit. **(b)** An important caveat is that the proof is formulated for a qubit (spin) system—i.e., a purely bosonic microscopic model with no explicit fermionic degrees of freedom, even if fermions can emerge effectively. If, instead, one introduces fermions at the microscopic level, for instance by considering an underlying Hubbard model rather than a purely spin model, these fermions interconvert the bosonic and fermionic toric codes and thereby unwind the finite-temperature topology. **(c)** The proof implies that there is a smoking gun measurement of fermionic statistics even in a finite temperature ensamble. It would be valuable to sharpen this into an explicit experimental proposal-even a Gedanken experiment-that makes concrete the formal exchange process illustrated in Fig. 2(b) and isolates the exchange statistics of the excitations, for example via an interferometric protocol in the spirit of Ref. [5]. [‡]

References

[1] Matthew B. Hastings, *Topological Order at Nonzero Temperature*, Phys. Rev. Lett. **107**, 210501 (2011).

[2] Tsung-Cheng Lu, Timothy H. Hsieh, and Tarun Grover, *Detecting Topological Order at Finite Temperature Using Entanglement Negativity*, Phys. Rev. Lett. **125**, 116801 (2020).

[3] Michael Levin and Xiao-Gang Wen, *Fermions, Strings, and Gauge Fields in Lattice Spin Models*, Phys. Rev. B **67**, 245316 (2003).

[4] Z. Li, D. Lee, and B. Yoshida, *How Much Entanglement Is Needed for Topological Codes and Mixed States with Anomalous Symmetry?* Phys. Rev. X **15**, 021090 (2025).

[5] Steven A. Kivelson and Chaitanya Murthy, *Modified Interferometer to Measure Anyonic Braiding Statistics*, Phys. Rev. Lett. **135**, 126605 (2025).

[‡]I thank Tarun Grover and Ruben Verresen for useful discussions.