Journal Club for Condensed Matter Physics DOI:10.36471/JCCM_October_2025_02
https://www.condmatjclub.org

Largest scale simulations

1. **Self-consistent tensor network method for correlated super-Moiré matter beyond one billion sites**
   Authors: Yitao Sun, Marcel Niedermeier, Tiago V. C. Antão, Adolfo O. Fumega, Jose L. Lado
   arXiv:2503.04373

2. **Solving the Gross-Pitaevskii equation on multiple different scales using the quantics tensor train representation**
   Authors: Marcel Niedermeier, Adrien Moulinas, Thibaud Louvet, Jose L. Lado, Xavier Waintal
   arXiv:2507.04262

Recommended with a Commentary by Anton Akhmerov, Kavli Institute of Nanoscience, Delft University of Technology

The title of the first paper intrigued me immediately. There is nothing particularly novel in Hartree–Fock mean-field simulations; all algorithms follow the same steps:

1. Choose a mean-field Hamiltonian
2. Compute the density matrix of all occupied states
3. Iterate the above until free energy is minimized (or self-consistency is achieved)

The second step—obtaining the density matrix—requires computing a high-rank dense operator. Such computations admit few optimizations. They cost $\mathcal{O}(N^3)$ operations for N orbitals and become highly impractical above $N \sim 10^4$ degrees of freedom on modern hardware. The paper, therefore, improves this result by an astounding amount: five orders of magnitude. The second paper simulates the time evolution of the Gross–Pitaevskii equation on a $10^6 \times 10^6$ grid. Although the Gross–Pitaevskii equation is not subject to the same cubic scaling as the Hartree–Fock approximation, storing such a grid would require roughly 12 terabytes of memory—far more than most users have available.

The main algorithmic idea that both papers use is the quantics tensor train (QTT) representation of the Hamiltonian, the wave function, and the density matrix. Long after matrix product states became a standard tool for representing 1D interacting quantum systems, Oseledets and others (see *e.g.* Ref. [1]) introduced the compactification of large arrays into tensor networks. After encoding matrix indices in binary, each index bit corresponds to a spin-1/2 degree of freedom: a matrix element becomes an amplitude for a particular configuration of these fictitious spins. Low-order (least significant) bits encode fine, local position information, while high-order (most significant) bits encode coarse-grained, large-scale position; correlations between different bit positions therefore map to correlations between different spatial scales. Put differently, entanglement between high-order bits captures long-wavelength (coarse) structure, whereas correlations among low-order bits describe short-range or high-frequency features. This is why a tensor-network approximation of the spin wavefunction naturally implements a multiscale compression of the original array, and why certain patterns of bit-to-bit coupling map onto correlations between ranges of Fourier modes in the original system. Constructing a tensor network approximation of this wave function provides an exponential compression of the original array, for a broad class of practically relevant operators.* The QTT representation is a relatively recent idea with applications to a variety of physical systems; after reading Ref. [2] I was tempted to rewrite our numerical codes using tensor networks.

Constructing the QTT representation of the tight-binding mean-field Hamiltonian is direct: the Hamiltonian terms are sparse, and QTT allows one to represent most position dependencies compactly. The real challenge, and the key idea of the first paper, is the construction of the density matrix. The authors use the kernel polynomial method [3]. Instead of applying Chebyshev polynomials of the Hamiltonian H to a small set of random vectors, they directly construct the QTT representation of each Chebyshev polynomial $T_n(H)$ using the recurrence relation $T_{n+1}(H) = 2HT_n(H) - T_{n-1}(H)$. The density matrix is then constructed approximately as a weighted sum of these polynomials, and the iteration to self-consistency follows the standard approach.

The main caveat of the approach becomes apparent when we revisit the problem and the physical phenomena we aim to capture. Specifically, for a simulation of $10^5 \times 10^5$ sites to be useful, something nontrivial must occur at that scale; otherwise, one could simulate a much smaller system and extrapolate the results, perhaps adding a smooth position dependence by interpolating the local solutions. Another subtlety lies in the limitations of the kernel polynomial method itself: computing the N-th Chebyshev polynomial of the Hamiltonian only allows one to capture correlations up to distances $\lesssim N$, and a much larger number of moments is required to reach sufficient precision for those correlations. Finally, constructing a fixed-rank QTT representation of $T_{n+1}(H)$ from $T_n(H)$ and $T_{n-1}(H)$ invariably introduces truncation errors; some correlations are neglected when the QTT bond dimension is fixed. The paper's analysis uses $N = 250$, which is much smaller than needed to capture relevant long-range physics. Finally, Hubbard interactions, together with the nearsightedness principle, typically lead to short-range correlations that are fully captured in smaller systems. I hope that the authors address these open questions in future works and demonstrate the performance of the method beyond proof of concept.

Despite these caveats, I find the algorithmic development worth very careful consideration. Firstly, even if there is a clear scale separation between complex local physics and emergent smooth long-range behavior, constructing a description that includes both—even approximately—is a daunting task, and it is precisely the problem the QTT approach ad-

*The vagueness of this formulation is intentional: I do not think we know what can and what cannot allow for a low-rank tensor approximation dresses. Secondly, a broad class of physics problems—such as the Gross–Pitaevskii equation studied in the second paper—naturally involve multiple scales, and the QTT approach appears well suited to them; it does not suffer from the error accumulation associated with sequential approximations. Even the question of approximating the density matrix in Hartree–Fock simulations may find a solution in the QTT representation: its strength is the ability to capture correlations across multiple scales. It may therefore be that alternative approaches to constructing the QTT density matrix would lead to improved results. Ultimately, while it is too early to draw firm conclusions, I find the application of tensor networks to large-scale simulations an exciting development worth following closely.

References

[1] TT-cross approximation for multidimensional arrays, I. Oseledets and E. Tyrtyshnikov, Linear Algebra Appl. **432**, 70-88 (2010).

[2] Learning tensor networks with tensor cross interpolation: New algorithms and libraries, Y. N. Fernández, M. K. Ritter, M. Jeannin, J.-W. Li, T. Kloss, T. Louvet, S. Terasaki, O. Parcollet, J. von Delft, H. Shinaoka, X. Waintal, SciPost Phys. **18**, 104 (2025).

[3] The kernel polynomial method, A. Weiße, G. Wellein, A. Alvermann, H. Fehske, Rev. Mod. Phys. **78**, 275 (2006).