

Implementing Quantum Cell Automata on a real neutral atom device

Rivan Rughubar, Tobias Stollenwerk
Forschungszentrum Jülich – PGI 12
Germany

Quantum cellular automata (QCA) provide a compelling framework for distributed quantum computation and the simulation of many-body quantum dynamics through the application of local, translation-invariant update rules. These update rules can also be used in the context of machine learning. By treating update rules as learnable functions, we are able to use repeated local interactions to learn complex global behaviour. There is an emerging body of work that applies QCA on neutral-atom (Rydberg) devices to solve a range of problems.

In this work we focus specifically on the maximum independent set (MIS) graph problem. MIS is a canonical NP-hard combinatorial optimization problem with applications ranging from network design and scheduling to error-correcting code construction. Solutions to the MIS problem can be found using a QCA approach, in which local update rules are applied iteratively to converge toward an optimal (or near-optimal) independent set.

Neutral-atom computers operate by trapping individual atoms in optical tweezers and selectively exciting them into Rydberg states. There are two features which make them well suited to the MIS problem. Firstly, the Rydberg blockade radius imposes a physically enforced, distance-dependent interaction that is inherently local, providing a hardware-native analogue of the locality constraint central to QCA update rules. Second, the flexible, reconfigurable geometry of tweezer arrays allows the qubit connectivity graph to be tailored to the structure of the problem instance.

Much of the existing research on QCA has been developed for idealized, theoretical neutral-atom devices. Real hardware introduces constraints that are not always considered in these theoretical treatments. Our goal is to map the QCA approach used for solving the MIS problem onto a real neutral-atom device. This involves mapping MIS graph instances onto the device's native hardware connectivity, as well as accounting for potential error correction. With this, we are able to discuss the current feasibility of QCA on near-term hardware and speculate on whether this approach may prove useful in a machine learning context.