

MAR14-2014-020812

Abstract for an Invited Paper
for the MAR14 Meeting of
the American Physical Society

Performance of simulated annealing, simulated quantum annealing and the D-Wave devices on hard spin glass instances

MATTHIAS TROYER, ETH Zurich

Quantum annealing - a finite temperature version of the quantum adiabatic algorithm - combines the classical technology of slow thermal cooling with quantum mechanical tunneling, to try bring a physical system towards its ground state. The Canadian company D-Wave systems has recently built and sold programmable devices that are designed to use this effect to find solutions to optimization problems. These devices raise many questions, in particular whether they outperform classical algorithms and show quantum speedup. To address this question I will start with an overview of simulated annealing (SA) and simulated quantum annealing (SQA). SA is a classical optimization method in which the cost function to be optimized is viewed as the energy of a physical system which is simulated in a Monte Carlo simulation. Slowly lowering the temperature in the simulation brings the system into a (local) energy minimum. SQA is a similar process, but where instead of lowering the temperature a control parameter is varied in quantum Monte Carlo (QMC) simulation, which (at constant temperature) brings a quantum system from the simple ground state of an initial Hamiltonian to a low-energy state of the final model implementing the same cost function. I will review evidence on the question whether SQA (or QA) is better at finding low energy states than classical SA. I will then present new results which show that SQA has problems in solving hard instances of spin glass problems: the distribution of times to solution exhibit very fat tails with slow power-law decays, indicating the presence of extremely hard to solve instances with the mean time to solution being dominated by rare instances. In SA, while there are also fat tails, they drop faster than in SQA, thus giving SA the edge when considering hard problem instances. Analyzing the performance of a 512-qubit D-Wave Two device we find the same issues with hard instances as in SQA and find no evidence for quantum speedup when comparing the performance on random spin glass instances to that of an optimized SA algorithm.