Simulated annealing versus quantum annealing

MATTHIAS TROYER, ETH Zurich

Based on simulated classical annealing and simulated quantum annealing using quantum Monte Carlo (QMC) simulations I will explore the question where physical or simulated quantum annealers may outperform classical optimization algorithms. Although the stochastic dynamics of QMC simulations is not the same as the unitary dynamics of a quantum system, I will first show that for the problem of quantum tunneling between two local minima both QMC simulations and a physical system exhibit the same scaling of tunneling times with barrier height. The scaling in both cases is $O(\Delta^2)$, where Δ is the tunneling splitting. An important consequence is that QMC simulations can be used to predict the performance of a quantum annealer for tunneling through a barrier. Furthermore, by using open instead of periodic boundary conditions in imaginary time, equivalent to a projector QMC algorithm, one obtains a quadratic speedup for QMC, and achieve linear scaling in Δ [1]. I will then address the apparent contradiction between experiments on a D-Wave 2 system that failed to see evidence of quantum speedup [2] and previous QMC results [3] that indicated an advantage of quantum annealing over classical annealing for spin glasses. We find that this contradiction is resolved by taking the continuous time limit in the QMC simulations which then agree with the experimentally observed behavior and show no speedup for 2D spin glasses. However, QMC simulations with large time steps gain further advantage: they "cheat by ignoring what happens during a (large) time step, and can thus outperform both simulated quantum annealers and classical annealers [4]. I will then address the question how to optimally run a simulated or physical quantum annealer. Investigating the behavior of the tails of the distribution of runtimes for very hard instances we find that adiabatically slow annealing is far from optimal. On the contrary, many repeated relatively fast annealing runs can be orders of magnitude faster for hard sin glass problems. The intuitive explanation is that hard instances, which are stuck in the wrong minimum can be solved faster by perturbing them [5]. I will finally discuss the consequences of these findings for designing better quantum annealers. [1] S.V. Isakov, G. Mazzola, V.N. Smelyanskiy, Z. Jiang, S. Boixo, H. Neven, and M. Troyer, arXiv:1510.08057. [2] T.F. Rønnow, Z. Wang, J. Job, S. Boixo, S.V. Isakov, D. Wecker, J.M. Martinis, D.A. Lidar, M. Troyer, Science **345**, 420 (2014). [3] G.E. Santoro, R. Martonak, E. Tosatti, and R. Car, Science **295**, 2427 (2002). [4] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, Science **348**, 215 (2015) (2015). [5] D.S. Steiger, T.F. Rønnow, M. Troyer, Phys. Rev. Lett. (in press); arXiv:1504.07991