MAR16-2015-020111

Abstract for an Invited Paper
for the MAR16 Meeting of
the American Physical Society

Quantum Annealing at Google: Recent Learnings and Next Steps

HARTMUT NEVEN, Google

Recently we studied optimization problems with rugged energy landscapes that featured tall and narrow energy barriers separating energy minima. We found that for a crafted problem of this kind, called the weak-strong cluster glass, the D-Wave 2X processor achieves a significant advantage in runtime scaling relative to Simulated Annealing (SA). For instances with 945 variables this results in a time-to-99%-success-probability 10^9 times shorter than SA running on a single core. When comparing to the Quantum Monte Carlo (QMC) algorithm we only observe a pre-factor advantage but the pre-factor is large, about 10^6 for an implementation on a single core. We should note that we expect QMC to scale like physical quantum annealing only for problems for which the tunneling transitions can be described by a dominant purely imaginary instanton. We expect these findings to carry over to other problems with similar energy landscapes. A class of practical interest are k-th order binary optimization problems. We studied 4-spin problems using numerical methods and found again that simulated quantum annealing has better scaling than SA. This leaves us with a final step to achieve a wall clock speedup of practical relevance. We need to develop an annealing architecture that supports embedding of k-th order binary optimization in a manner that preserves the runtime advantage seen prior to embedding.