

Recent developments in quantum annealing

*Iliia Zintchenko, Ethan Brown, and Matthias Troyer
Theoretical Physics, ETH Zurich, 8093 Zurich, Switzerland
December 7, 2015*

Summary

We review and comment on some of the recent developments in quantum annealing. In particular, we will comment on recent findings that quantum annealers may be best suited for finding a class of approximate solutions and on Google's announcement of quantum annealers outperforming classical annealers by significant margins. This is not meant to be a comprehensive review and we apologize in advance to all authors whose relevant papers we do not mention.

Introduction

Quantum annealing is an optimization method that employs quantum effects to escape local minima of the cost function by “tunneling” through barriers separating local minima. If an optimization problem exhibits very rough cost functions with tall but narrow barriers this can be more efficient than classical annealing which climbs over barriers.

Quantum annealing has been commercialized in the devices of D-Wave Systems (Burnaby, BC, Canada), which are analog quantum devices designed to use quantum annealing to solve an NP-hard binary quadratic optimization problem with cost function $\sum_{i,j} a_{ij}x_i x_j + \sum_i b_i x_i$ and variables x_i that can take values 0 or 1. The latest D-Wave device, D-Wave 2X, can solve optimization problems with up to 1152 variables, with couplings a_{ij} that are nonzero on a certain sparse graph.

While quantum devices will not solve NP-complete problems in polynomial time, they may still have a (potentially large) advantage over classical algorithms. However, in contrast to certain algorithms for universal quantum computers that are known to have exponential speedup over the best known classical algorithms, it is still unknown whether and where analog quantum annealers have a computational advantage over classical algorithms.

Performance of D-Wave and the sweet spot for quantum annealing

Experiments on the D-Wave devices over the past years have indicated that the devices use quantum effects [1,2], and that their behavior is consistent with quantum annealing [3]. However, it has also been shown that for the class of problems used in initial benchmarks, the performance of the devices also correlates well with a classical (mean-field) approximation of quantum annealing [4], which raises questions about whether the devices can outperform classical optimizers. While first benchmarks comparing to commercial optimization algorithms showed a 3600x advantage for D-Wave 2, subsequent tests comparing to optimized state-of-the-art classical codes [5] showed performance comparable to that of a high-end Intel Xeon CPU. More importantly no scaling advantage (“quantum speedup”) has so far been seen.

Given unavoidable calibration inaccuracies of any analog device, one may focus on using quantum annealers to find good, but not necessarily optimal solutions. Indeed, both experiments by D-Wave [6] and simulations [7] have shown that quantum annealers are efficient in quickly finding good solutions– but then are slow in further improving the

quality of the solution. This behavior can be explained by realizing that a quantum annealer can very quickly go to lower and lower energies as long as barriers are narrow, but ultimately gets stuck once broad barriers are encountered [8].

The sweet spot for quantum annealers is thus quickly finding approximate solutions. Using the energies that D-Wave devices can find easily as the “target” in a benchmark aimed at showcasing the performance of the devices, D-Wave has performed a careful study showing that their latest device is 15x faster than the best optimized classical codes on a single core of an Intel CPU (and thus still comparable to a high-end multi-core CPU) in finding good but not necessarily optimal solutions [6].

How useful a quantum annealer is at finding approximate solutions will depend on whether the quality of the found solutions is good enough for a specific application and how hard it is to find similar approximations using approximate classical solvers.

Significant quantum effects in annealing

Google has recently presented evidence for significant quantum effects on the latest D-Wave 2X device using problem instances that are designed to be hard for thermal annealing algorithms but easy for quantum annealers [8]. These instances strongly couple groups of eight variables, thus enforcing them to have the same values. This makes it hard to change all variables in this group by methods that change one at a time, since large energy penalties have to be paid at intermediate steps while changing one variable after the other. A quantum annealer, in contrast, can collectively change all variables at once.

Google’s results indeed show a huge advantage of D-Wave 2X over simulated classical annealing on these carefully chosen native instances, demonstrating that D-Wave 2X uses quantum effects and showing that quantum effects can be an important resource in annealing. These results are a significant proof-of-principle demonstration that quantum effects in annealing can, in some cases, lead to substantial performance gains over purely classical annealing.

However, this should not be taken as a demonstration of quantum annealers outperforming classical optimization methods. Already a small variation of simulated annealing [9] significantly speeds up annealing and reduces the time to solve these problems from months (as reported in [8]) to under a second. The Google team indeed mentions that even more efficient classical optimization algorithms exist for these problems, which outperform the current D-Wave 2X device for most problem instances. Furthermore, these results are for native instances that perfectly fit the hardware graph of the device. For generic problems that do not map well onto the hardware of a quantum annealer, performance will suffer significantly.

It will now be important to further explore if there are problems where quantum annealing has advantages over even the best classical algorithms, and to find if there are classes of application problems where such advantages can be realized.

The potential of quantum annealing

The performance results reported for D-Wave 2X are not the final answer. With improved engineering, especially faster annealing and readout we expect that the time to perform a quantum annealing run can be reduced by a factor 100x over the current generation devices.

However, there are also significant challenges that will hamper the performance of analog quantum annealers. In particular, there are problems due to *misspecification* of

the cost function due to *calibration inaccuracies*. These restrict the range of applications to those where the terms in the cost function do not differ widely. Furthermore, they limit the quality of solution that can be reached and increasingly better calibration will be required as one scales to larger systems. Also, often efficient classical approximation schemes exist that can efficiently reach the same quality of solution.

Another challenge is the *embedding* of problems into the native hardware architecture with *limited connectivity*. Embedding a generic quadratic binary optimization problem with N variables requires N^2 qubits and couplers, representing one variable by up to N qubits. Such embeddings are furthermore expected to significantly slow down the annealing dynamics. More complex optimization problems may incur even higher mapping and embedding penalties.

Finally, there is the open question of quantum speedup in analog quantum annealing. For a broad class of tunneling problems classical algorithms simulating quantum systems (through quantum Monte Carlo simulations) can profit from quantum tunneling through barriers in the same way as stoquastic quantum annealers (such as the D-Wave devices) [10]. New types of non-stoquastic quantum annealers may thus be needed to get more than a constant advantage over classical devices.

These challenges will need to be addressed in the development of future improved quantum annealers. For this reason, significantly longer coherence times, higher calibration accuracy and range of couplings, better connectivity of the hardware, and non-stoquastic quantum dynamics are critical components of the quantum enhanced optimization (QEO) efforts of the Intelligence Advanced Research Projects Agency (IARPA) of the U.S. government. In the longer term future, these issues will also be solved by implementing quantum annealing on future circuit-model universal quantum computers that will be more flexible both in terms of connectivity and choice of cost function.

References

- [1] T. Lanting *et al.*, *Entanglement in a Quantum Annealing Processor*, Phys. Rev. X **4**, 021041 (2015)
- [2] Sergio Boixo *et al.*, *Computational Role of Multiqubit Tunneling in a Quantum Annealer*, arXiv:1502.05754
- [3] S. Boixo *et al.*, *Evidence for Quantum annealing on more than a hundred qubits*, Nature Physics **10**, 218 (2014)
- [4] S.W. Shin, G. Smith, J. A. Smolin, U. Vazirani, *How "Quantum" is the D-Wave Machine?*, arXiv:1401.7087
- [5] S.V. Isakov, I.N. Zintchenko, T.F. Rønnow, M. Troyer, *Optimized simulated annealing for Ising spin glasses*, Computer Physics Communications, **192**, 265 (2015)
- [6] J. King, S. Yarkoni, M. M. Nevisi, J. P. Hilton, C. C. McGeoch, *Benchmarking a quantum annealing processor with the time-to-target metric*, arXiv:1508.05087
- [7] B. Heim, T. F. Rønnow, D. V. Isakov, and M. Troyer, *Quantum versus Classical Annealing of Ising Spin Glasses*, Science **348**, 215 (2015)
- [8] V. S. Denchev *et al.*, *What is the Computational Value of Finite Range Tunneling?*, . arXiv:1512.02206
- [9] I. Zintchenko, M. B. Hastings, M. Troyer, *From local to global ground states in Ising spin glasses*, Phys. Rev. B **91**, 024201, 2015
- [10] S.V. Isakov *et al.*, *Understanding Quantum Tunneling through Quantum Monte Carlo Simulations*, arXiv:1510.08057