The Quantum Computing Company™

Quantum Performance Evaluation: A Short Reading List

WHITEPAPER

Summary

This white paper presents a brief survey of recent work on empirical performance evaluation of D-Wave systems.

1 Introduction

Since the first release of D-Wave annealing-based quantum computers in 2010, scores of research papers have been published describing their physical properties, capabilities, and performance. The research domain is complex and rich, which means that the work is ongoing and will continue for many years.

This white paper gives a snapshot of recent work on quantum system performance evaluation, which considers both solution quality and computation time. This work may be roughly divided into four categories:

- **Benchmarking:** evaluating quantum processor performance at solving application problems.
- **Tuning**: developing techniques and tools for improving performance.
- **Mapping the performance landscape**: characterizing input properties that affect solution quality and/or computation time.
- **Fundamental properties:** addressing open questions motivated by theoretical work.

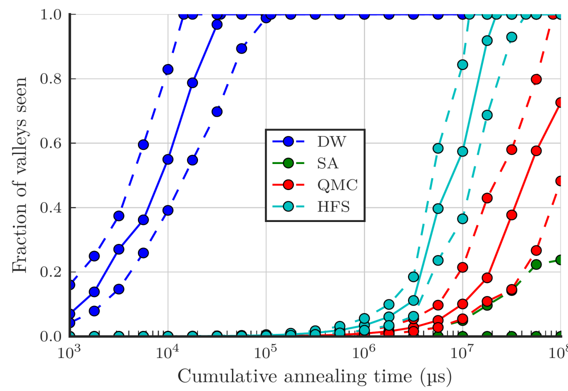

Figure 1: In this example [3], the D-Wave 2000Q quantum processing unit (QPU) finds samples of optimal solutions about 1000 times faster than three classical alternatives.

2 Applications Benchmarking

Although scores of papers have been published showing how to translate application problems for solution on D-Wave systems, relatively few have discussed performance. This is primarily because previous-generation processors have been too small to hold problems of interesting size: when problems are small, classical and quantum solvers are uniformly fast, and unequivocal performance distinctions cannot be made.

Two recent papers overcome that obstacle to demonstrate the potential for significant performance advantages as qubit counts continue to grow:

- Trummer and Koch [6] look at a problem in database optimization, using realistic input structures that maximize the number of variables packed onto the quantum processor. They report that the best of five classical solvers can be up to 1000 times slower at finding solutions of comparable quality to those found by a D-Wave 2X system.

- Ushijima-Mwesigwa et al. [7] look at a problem motivated by simulations of quantum molecular dynamics, which involves finding good decompositions of large density matrices into subsystems. They report that quantum and hybrid classical–quantum approaches implemented on current systems can equal or outperform state of the art classical methods.

3 Performance Tuning

Another body of work aims to identify techniques for making best use of quantum processor capabilities. Here are three examples.

- Pudenz et al. [5] show how adding a small percentage of error-control qubits to problems formulated for D-Wave processors can improve success probabilities.
- Bian et al. [2] describe application of a D-Wave 2000Q system to a problem in constraint programming and circuit fault detection. They show how choice of strategies for problem transformation and embedding can affect processor performance.
- Andriyash et al. [1] show how to improve performance at integer factoring by modifying the shape of the quantum annealing path.

4 Mapping the Performance Landscape

A growing family of input properties — with labels such as frustration, glassiness, and floppiness — have been identified that can be used to predict quantum as well as classical solver performance. King et al. [4] look at two properties, *local ruggedness* and *global frustration*, that can be challenging for classical approaches to solve efficiently, but are not especially hard for quantum annealing. Tests on inputs having these properties show that anneal times for a D-Wave 2000Q processor can be about three orders of magnitude faster at finding a single optimal solution, and at sampling all optimal solutions (see Fig. 1).

5 Fundamental Properties

Prompted by open questions in the theoretical foundations of quantum computing, several researchers have focused on questions of validating, quantifying, and characterizing the quantum mechanical effects in D-Wave technologies. A related thread of research looks at a property known as *quantum speedup* that involves comparison of how measures of classical and quantum work scale with respect to problem size and hardness.

Job and Lidar [3] present a comprehensive survey of empirical work on quantum validation, quantum speedup, and quantum error correction.

References

[1] E. Andriyash et al., Boosting integer factoring performance via quantum annealing offsets, D-Wave TR 14-1002A-B, 2016.

[2] Z. Bian et al., Mapping constrained optimization problems to quantum annealing with application to fault diagnosis, *Frontiers in ICT*, 2016.

[3] J. Job and D. Lidar, Test driving 1000 qubits, arXiv 1706.07124, 2017.

[4] J. King et al., Quantum annealing amid local ruggedness and global frustration, D-Wave TR 14-1003-C, 2017.

[5] K. L. Pudenz et al., Error corrected quantum annealing with hundreds of qubits, *Nature Communications* 5, 3243, 2014.

[6] I. Trummer and C. Koch, Multiple query optimization on the D-Wave 2X adiabatic quantum computer, *VLDB* 2016.

[7] H. Ushijima-Mwesigwa et al., Graph partitioning using quantum annealing on the D-Wave system. arXiv 1705.03082v1, 2017.