



Research | December 01, 2017

# How Will Early Quantum Computing Benefit Computational Methods?

By Phil Goddard, Susan Mniszewski, Florian Neukart, Scott Pakin, and Steve Reinhardt

The first half of 2017 saw considerable news about the imminent impact of quantum computing (QC) on science, math, and data analysis [9]. Diverse viewpoints predict wildly-varying timescales—from years to decades—for initial impact, and debate what types of problems will first be affected. At the 2017 SIAM Annual Meeting, held in Pittsburgh, Pa., this July, we organized a minisymposium entitled “Identifying Computational Methods for Early Benefit from Quantum Computing,” which addressed how QC will become useful to practicing computational scientists. We aimed to identify computational methods that are likely to benefit from QC in the next five years, and share the experiences of early QC application developers.

The conceptual appeal of QC systems is the use of their computational power, which grows exponentially with the number of quantum bits (or *qubits*), to reduce the time spent executing combinatorial algorithms, which on classical computers grows exponentially with the number of variables. The most advanced commercially available quantum computers are currently delivered by D-Wave Systems; the number of qubits in these systems has doubled every other year for the past six years, a trend that is expected to continue. We believe that D-Wave quantum computers will deliver high-quality solutions quickly and in a sustainable way.

## Quantum Annealing

Scott Pakin of Los Alamos National Laboratory (LANL) described the current state of QC systems, where contemporary quantum computers implement one of two computational models: the *gate* or *circuit* model [4], or the *quantum annealing* (QA) model [5]. Here we focus on QA-based quantum computers, as they are the most commercially advanced.

A QA-based processor is a special-purpose device that natively solves the *quadratic unconstrained binary optimization* (QUBO) problem. The goal of a QUBO is to find the  $x_i \in \{0, 1\}$  that minimize

$$x^T Q x \quad (1)$$

for upper-triangular matrix  $Q$  with  $Q_{i,j} \in \mathbb{R}$ . The QUBO form is equivalent to the *unconstrained binary quadratic programming* (UBQP) problem and the Ising model.

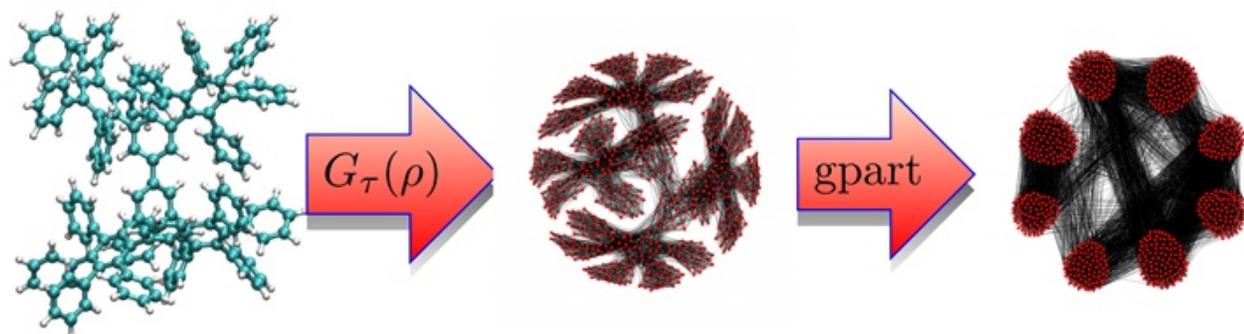
One can think of a quantum annealer as a hardware implementation of simulated annealing that uses quantum effects—superpositioning, entanglement, and quantum tunneling—to reduce the likelihood of getting stuck in a local minimum. In both quantum and simulated annealers, finding the  $x_i$  that truly minimize (1) is not guaranteed. Instead, the goal is normally to find good solutions quickly. A D-Wave quantum computer can propose a solution to (1) in only  $5\mu\text{s}$  even for  $N > 2000$  (where the search space is greater than  $2^{2000}$ ), as supported by a D-Wave 2000Q™ system.

Mapping computational problems into QUBO form, which requires creativity and ingenuity but can lead to encouraging results when accomplished, is a key challenge that all subsequently-described applications face.

## Applications

**Graph Partitioning.** Combinatorial optimization as graph-theoretic problems is ubiquitous throughout mathematics, computer science, physics, chemistry, bioscience, machine learning, and complex systems. Many of these problems are NP-hard and therefore rely on heuristic solutions. The availability of quantum computers provides new opportunities to explore relevant quantum graph algorithms, particularly novel QA heuristics and hybrid methods that may out-perform classical approaches or efficiently solve new problem types.

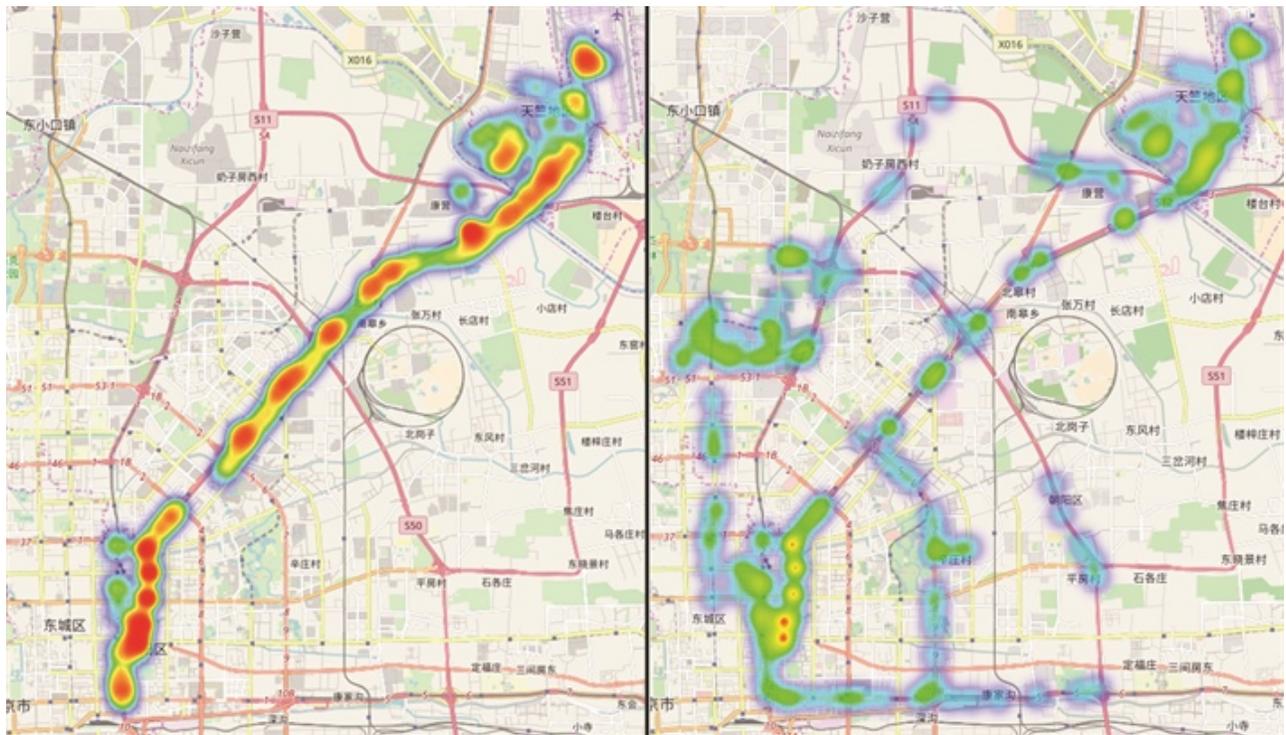
Sue Mniszewski of LANL presented her team’s work on graph partitioning algorithms using QA on the D-Wave 2X™ system [6], motivated by recent studies on graph-based electronic structure theory. Graph partitioning is a natural fit to the QUBO form solved by the D-Wave 2X system. One can naturally map unconstrained community clustering—based on the modularity metric—onto the QUBO, despite its status as a “very hard” problem for a classical computer. When constraints are imposed (e.g., equal partitioning), QUBO reformulation is necessary. Mniszewski demonstrated concurrent partitioning based on the graph coloring problem, community clustering via modularity, and iterative multilevel partitioning with refinement (similar to METIS and KaHIP) on benchmark graphs and electronic structure graphs. Even graphs with approximately self-similar structure, which are traditionally difficult to partition, benefit from quantum graph partitioning (see Figure 1). Graph partitioning results from quantum and hybrid classical-quantum approaches show that the D-Wave 2X system is comparable to current state-of-the-art methods, and sometimes better.



**Figure 1.** From left to right: Phenyl dendrimer molecular structure to the electronic structure graph and the resulting 8-concurrent graph partitioning using a D-Wave system. Image credit: Susan Mniszewski, Los Alamos National Laboratory.

**Traffic Flow Optimization.** Florian Neukart from Volkswagen (VW) described the mapping of certain parts of the real-world problem of maximizing traffic flow into a form suitable for QA [7]. The objective is to reduce the time required for a set of cars to travel between their individual sources and destinations by minimizing total congestion over all road segments. To combat the limited size and connectivity of current-generation D-Wave systems, the VW team developed a hybrid quantum and classical workflow—based on the T-Drive trajectory dataset of the cars' GPS coordinates—that mimics traffic flow optimization in near-real time, integrating calls to the D-Wave system.

The team maximized simulated traffic flow between the Beijing city center and airport by redirecting a subset of 418 selected taxis to alternate routes, thus minimizing the number of *congested* (i.e., having more than a given number of vehicles at a given time) road segments. This required simultaneous optimization over all vehicles. The program considered three alternative routes, described by lists of street segments, for each trip between the city center and airport, and assigned cars to routes that minimized the number of congested road segments. The left half of Figure 2 displays the traffic density on the Beijing road graph before optimization, and the right displays the density after optimization via D-Wave's qbsolv tool [2].



**Figure 2.** Heatmaps of traffic congestion in unoptimized (left) and optimized (right) flows. Image credit: Florian Neukart, Volkswagen Group of America.

The VW team focused on finding good quality solutions within short periods of calculation. They ultimately showed that time-critical optimization tasks, such as continuous redistribution of position data for cars in dense road networks, are suitable candidates for quantum applications.

**Finance.** Phil Goddard of 1QB Information Technologies (1QBit) works with applications in computational finance, where many high-value problems can be cast as combinatorial optimization problems and are hence solvable with QA. Specific applications include multiperiod mean-variance portfolio optimization, (quantum) hierarchical risk parity feature selection for credit scoring, and tax loss harvesting [1]. While adhering to realistic investment constraints, the problem variables can be discretized and reduced through suitable transformations to the QUBO form solved by a D-Wave system. The size of the solvable optimization problem (measured in number of problem variables) depends on several factors, including the technique used to discretize the problem and the way it is embedded onto the Chimera graph topology in which the D-Wave system's qubits are laid out. Conceptually, the benefits of using QC will increase with problem size.

## Programming

Steve Reinhardt of D-Wave Systems described the range of software tools available to solve problems on existing commercial quantum computers. Only in the last year or two have interfaces emerged that are higher level than the system-dependent quantum machine instruction exposed by D-Wave's low-level system interface [3]. Several groups of tool developers deliver interfaces that allow a problem to be stated as a polynomial and then converted to a QUBO. Two groups deliver *solvers* that solve a problem larger than that natively supported by the hardware by repeatedly decomposing the problem into chunks that can be executed on the system. These tools have become essential for application and higher-level tool development. QC Ware Platform [10] and edif2qasm [8] are tools that enable subject-matter experts to express problems in high-level forms and map those problems to a QUBO for D-Wave execution.

## Summary

When comparing techniques for the different applications, we see that effectively mapping a problem to QUBO form is essential to obtaining good results. Some problems map directly to QUBOs while others map with intermediate steps that may add ancillary variables. Either approach matches well with a D-Wave quantum computer, ideally resulting in a problem that is sufficiently hard, as D-Wave systems are more likely to deliver a performance advantage over classical systems in the case of hard problems. Additionally, programs solving real-world problems invariably execute in a hybrid classical-quantum mode. All of the aforementioned applications also take advantage of a decomposing solver to fit the problem—in chunks—onto the QC system.

### References

[1] 1QB Information Technologies. *1Qbit: Our Thinking*. Retrieved from <https://1qbit.com/our-thinking/>.

[2] Booth, M., Reinhardt, S.P., & Roy, A. (2017). *Partitioning Optimization Problems for Hybrid Classical/Quantum Execution*

- (D-Wave Technical Report Series, 14-1006A-A). Retrieved from [https://github.com/dwavesystems/qbsolv/blob/master/qbsolv\\_techReport.pdf](https://github.com/dwavesystems/qbsolv/blob/master/qbsolv_techReport.pdf).
- [3] *Developer Guide for C*. (2017). Burnaby, British Columbia, Canada: D-Wave Systems, Inc.
- [4] Feynman, R.P. (1986). Quantum mechanical computers. *Found. Phys.*, 16(6), 507-531.
- [5] Kadowaki, T., & Nishimori, H. (1998). Quantum annealing in the transverse Ising model. *Phys. Rev. E*, 58, 5355-5363.
- [6] Mniszewski, S.M., Ushijima-Mwesigwa, H., & Negre, C.F.A. (2017). *Graph partitioning using the D-Wave for electronic structure problems*. In *Identifying Computational Methods for Early Benefit from Quantum Computing* (minisymposium). 2017 SIAM Annual Meeting. Pittsburgh, PA, July 10-14, 2017.
- [7] Neukart, F., Von Dollen, D., Compostella, G., Seidel, C., Yarkoni, S., & Parney, B. (2017). Traffic flow optimization using a quantum annealer. Preprint, *arXiv:1708.01625*.
- [8] Pakin, S. (2017). *edif2qasm: Run Hardware Descriptions on a Quantum Annealer*. Retrieved from <https://github.com/lanl/edif2qasm>.
- [9] Palmer, J. (2017). Here, There, and Everywhere: Quantum technology is beginning to come into its own. *The Economist, Technology Quarterly*. Retrieved from <http://www.economist.com/news/essays/21717782-quantum-technology-beginning-come-its-own>.
- [10] QC Ware Platform software. (2017). *QC Ware*. Retrieved from <https://platform.qcware.com/>.

Phil Goddard is Head of Research at 1QBit, and president and principal consultant at Goddard Consulting. He specializes in applying cutting-edge techniques to high-value optimization and machine learning problems. Susan Mniszewski is a computer scientist at Los Alamos National Laboratory (LANL). Her research interests include solving optimization problems using quantum computing. Florian Neukart works as a principal scientist in the fields of artificial intelligence and quantum computing at Volkswagen Group of America, where he is conducting research on self-driving vehicles, quantum machine learning, self-optimizing robots, and artificial intelligence. Scott Pakin is a computer scientist who currently serves as the technical/scientific point of contact for LANL's D-Wave quantum annealer. He has researched numerous areas of high-performance computing. Steve Reinhardt has built hardware/software systems that deliver new levels of performance via conceptually simple interfaces. He now leads D-Wave's efforts with customers to map early applications to D-Wave systems.