

Benchmarking quantum annealing for complex network analysis on synthetic social networks

Ojas Parekh
Sandia National Labs

A question?

Which is harder (pick a reasonable definition of “hard”)?

Finding an optimal solution
99% of the time

Finding a solution within
99% of optimal all the time

Which is more practically relevant?

A question?

Which is harder (pick a reasonable definition of “hard”)?

Finding an optimal solution
99% of the time

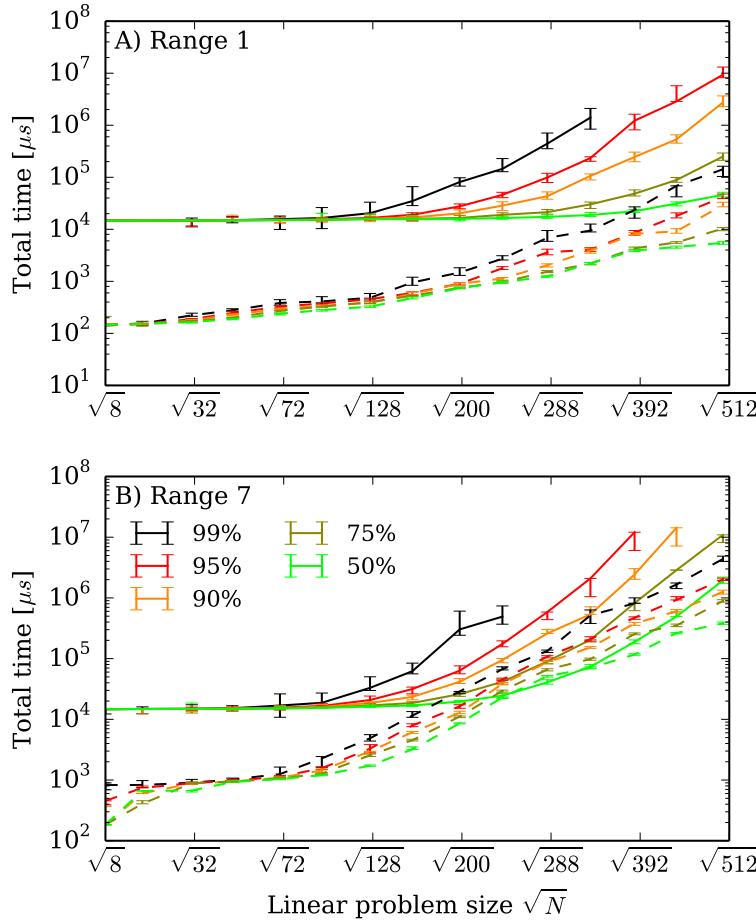
Finding a solution within
99% of optimal all the time

“Average”-case analysis,
rather than worst-case

Approximation algorithm,
or approximation scheme

How should we measure success?

Optimal 99% of the time

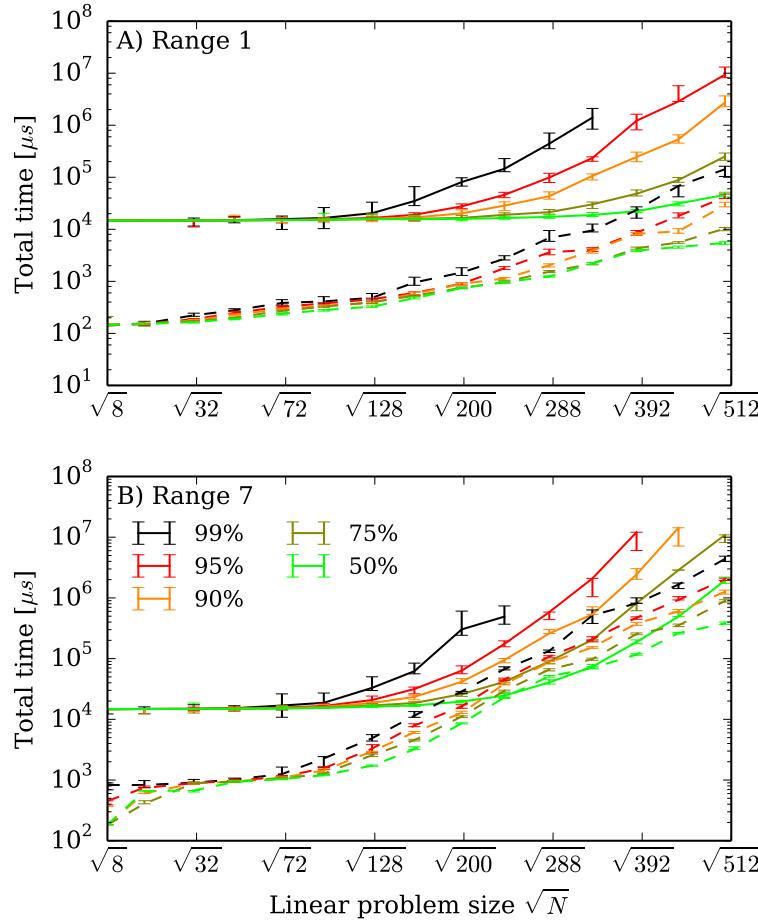


Instances: D-Wave 2 Chimera graph
with randomly chosen coupler weights

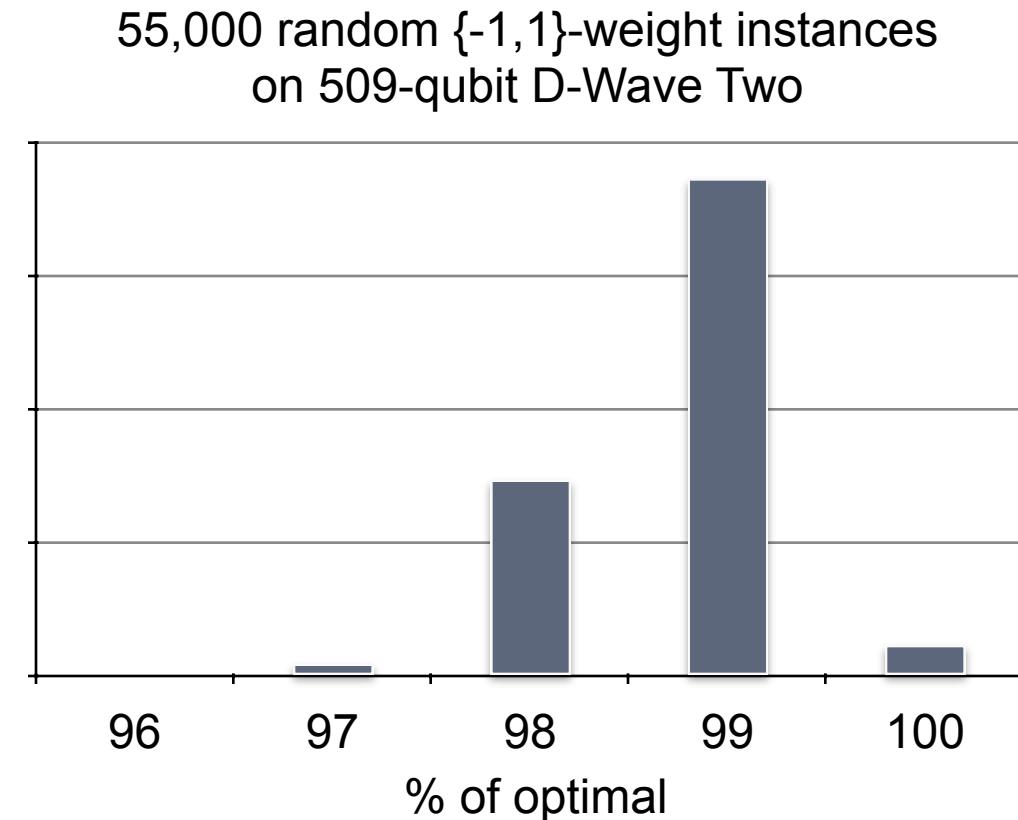
“Exponential” scaling as $c^{\sqrt{n}}$ for both
D-Wave 2 and simulated annealing

How should we measure success?

Optimal 99% of the time



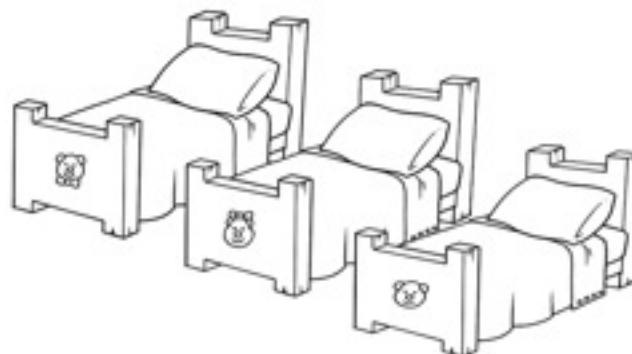
Within 99% of optimal all of the time



Was always within 96% of optimal!

We ask

- What is an appropriate measure of success?
 - Which are of practical significance?
 - Which lead to good science?
- What classical algorithm(s) should be used for comparison?
 - How should they be configured?
 - Fair comparison of classical and quantum resources?
- How should one select appropriate benchmark instances?
 - Are purely random instances meaningful?
 - How do we bridge the gap to complex real-world instances?

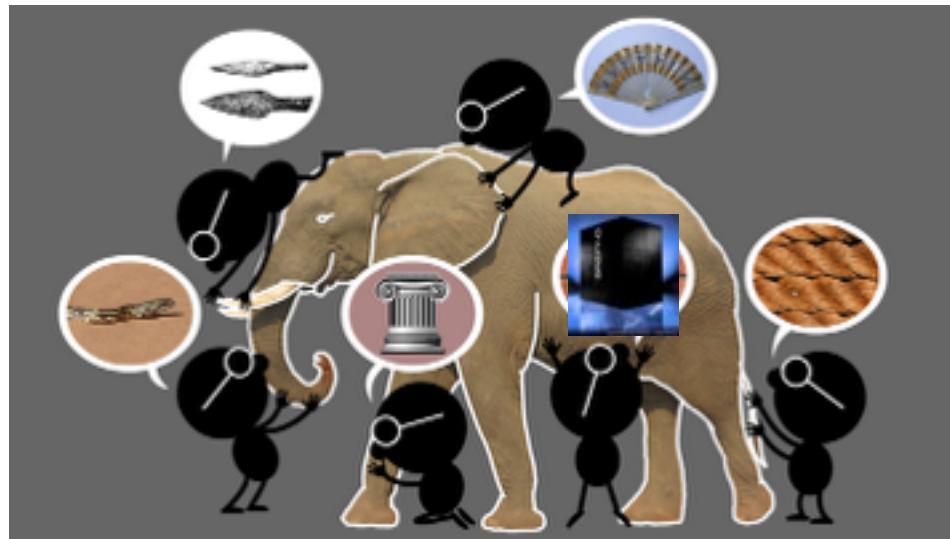


Benchmarking goals

- Assess current and projected future performance of D-Wave
- Compare against variety of classical algorithms, including real-world heuristics on real-world-like instances
- Identify barriers and work-arounds in formulating and representing problems on AQCs
- Place D-Wave community benchmarking results in context
- Gap between quantum algorithms and quantum software
 - (Theoretical) algorithms: worst-case, big-O running times and P vs NP issue
 - (Real-world) software
 - Instances of NP-hard problems solved routinely
 - Software engineering and factors hidden by big-O matter
 - New field: *Quantum Software Engineering*
 - Quantum software for specific architectures rather than pseudo-code
 - Algorithms must be designed to leverage architectural limitations

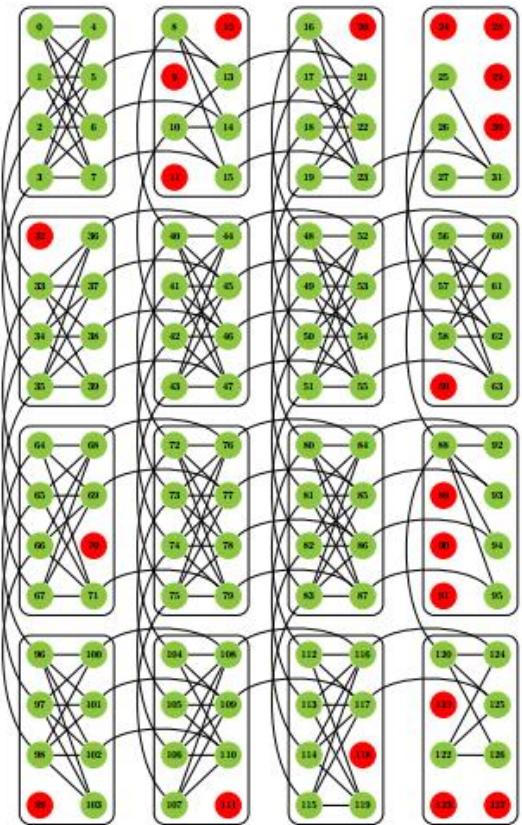
The Elephant in the room

- Sandia focus: general D-Wave-agnostic (A)QC
- D-Wave is a specific technology (superconducting flux qubit) and architecture
- We ask more general questions: how do quantum resources and architectural implementations affect problem-solving capacity?
- Other qubit technologies at Sandia: silicon, trapped neutral atoms, and trapped ions

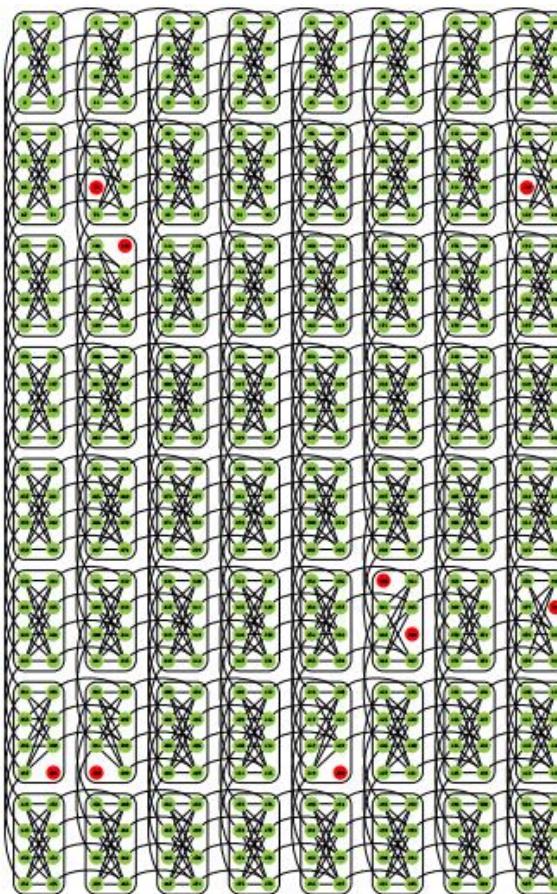


Comparison of Rainier and Vesuvius chips

Rainier
108/128
spins



D-Wave One

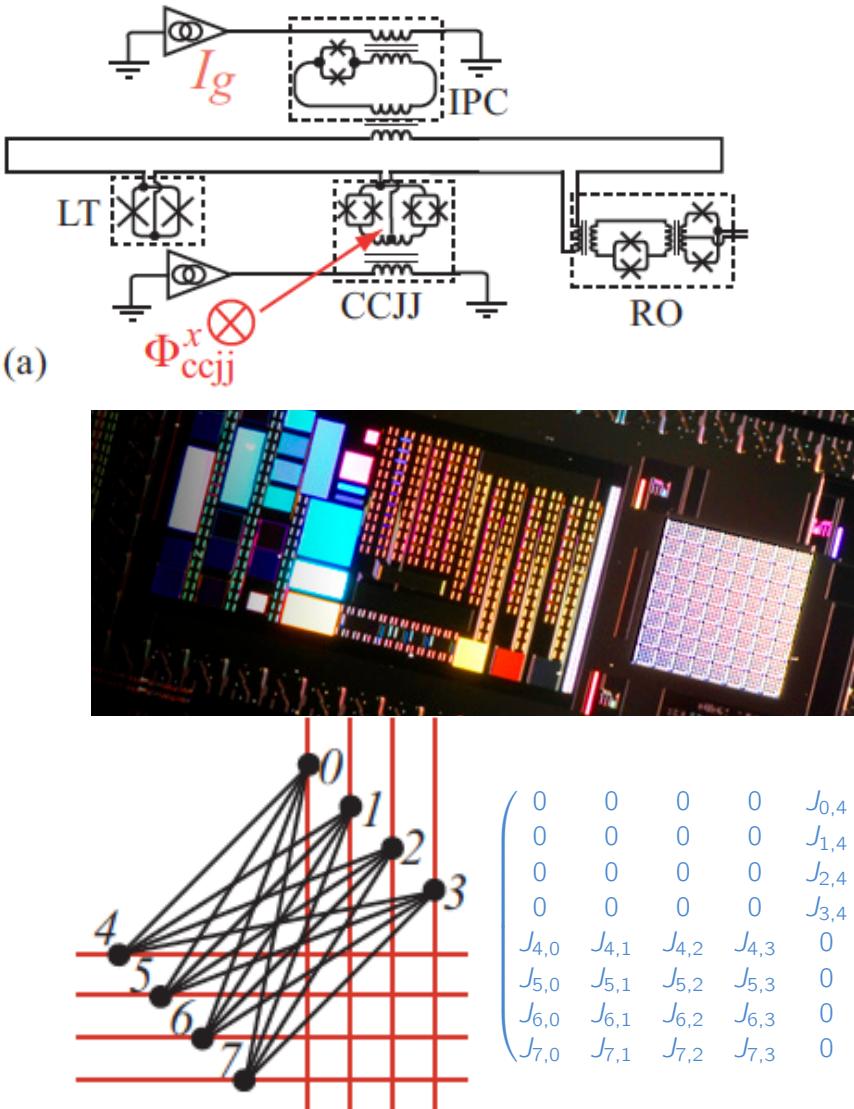
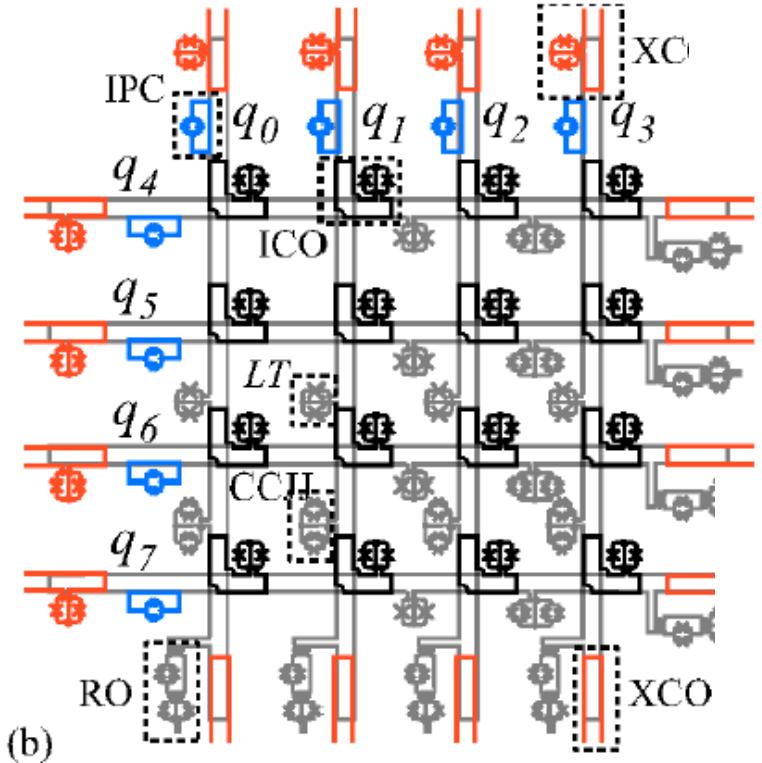


D-Wave Two

Vesuvius
506/512
spins

Images from D-Wave Systems: <http://www.dwavesys.com> .

Eight qubit cell architecture



R. Harris et al., Phys. Rev. B 82, 024511 (2010)

QUBO

(Quadratic Unconstrained Binary Optimization)

$$\text{QUBO: } f(x) = \min_{x \in \{0,1\}^n} b^T x + x^T C x$$

- Well-suited for discrete optimization applications
- Variables correspond to qubits
More qubits = richer problem modeling
- Matrix C corresponds to a graph
- AQC architectural constraints and hardware dictate edges and weight ranges

$$(b, C) \rightarrow$$

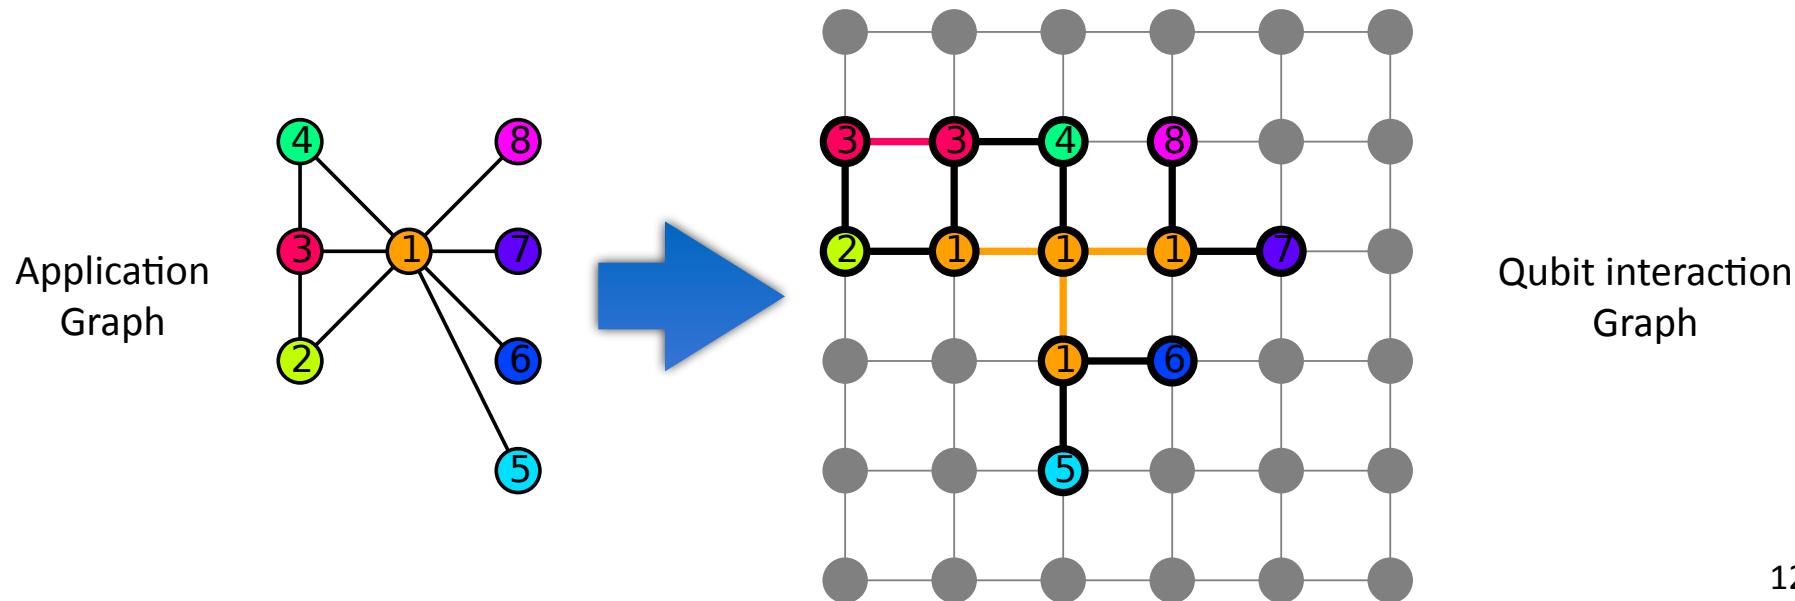
QUBO
AQC

$$\rightarrow x$$

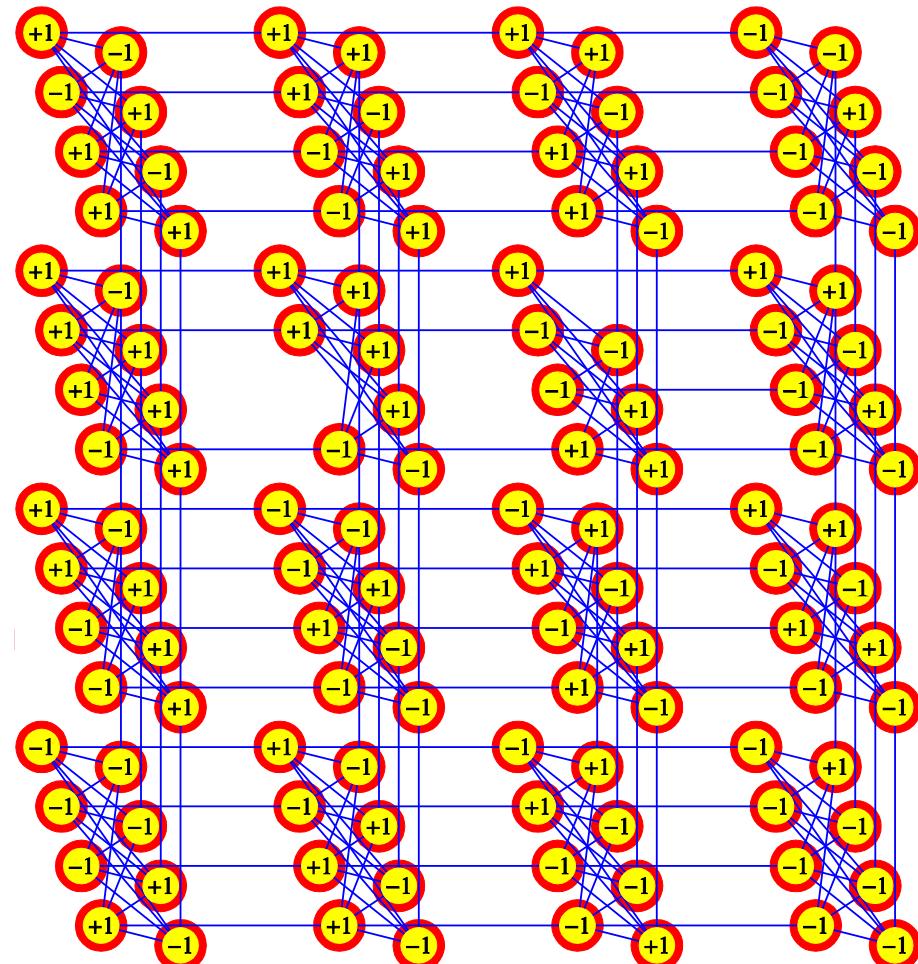
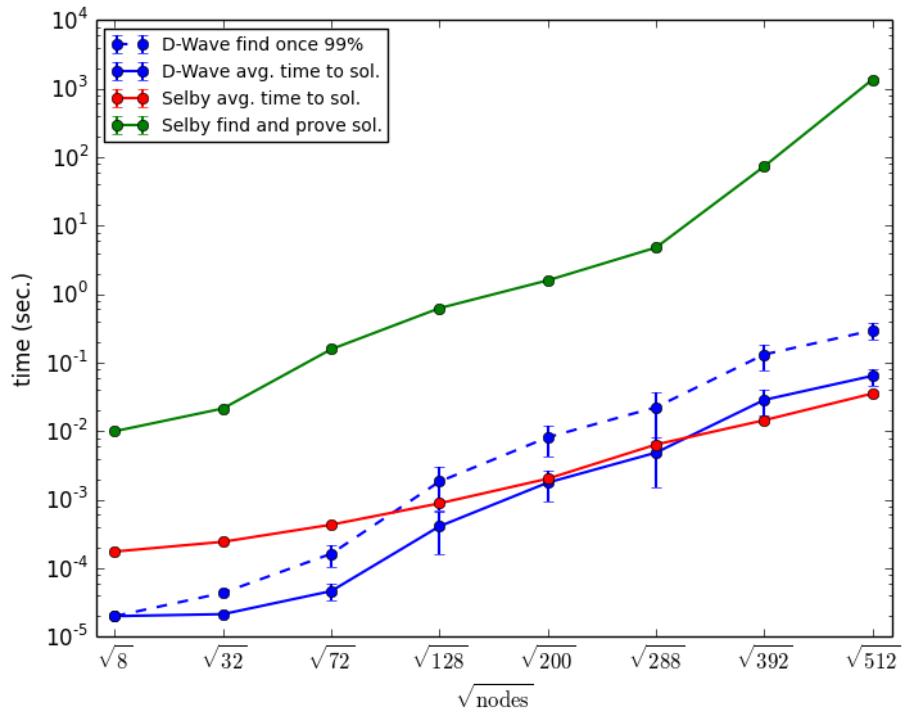

Sandia
National
Laboratories

Solving problems with D-Wave: challenges

- Application graph must be embedded within Chimera graph
 - Requires extra qubits; worst case: $O(n^2)$ qubits for n nodes
 - Very hard to determine a good embedding for a given graph
 - Typical approaches to embedding require large weights to force all qubits corresponding to a node have same spin
 - Efficiency: #(application graph nodes)/#(qubits in Chimera graph)

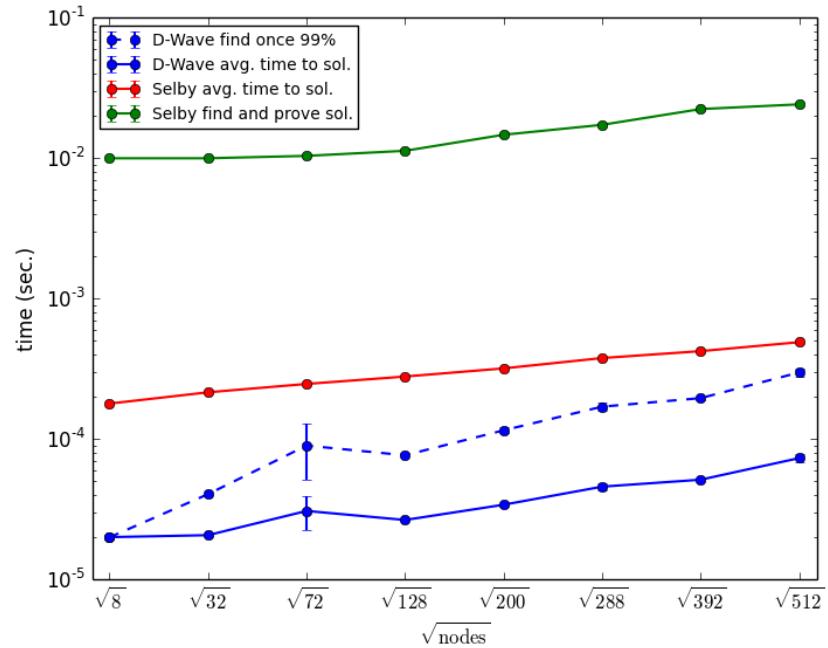
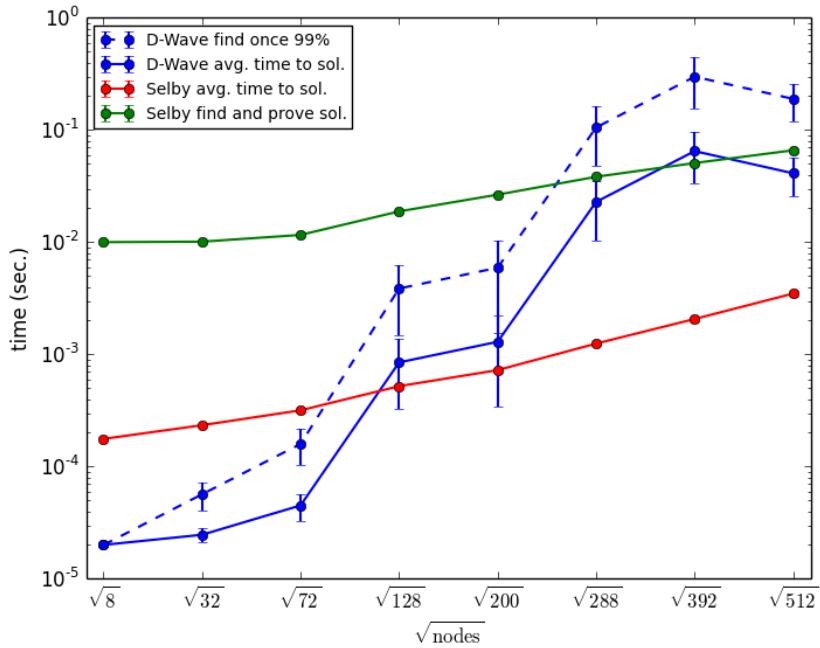


Example: hard problem

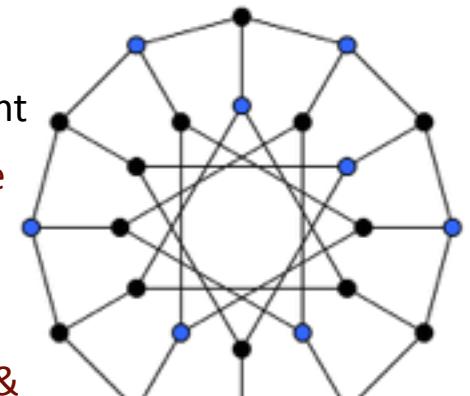


- Random Ising instances on Chimera graph
- Pro: avoids embedding
- Con: limited practical significance
- Standard instances for D-Wave benchmarking
- We proved that these instances are NP-hard

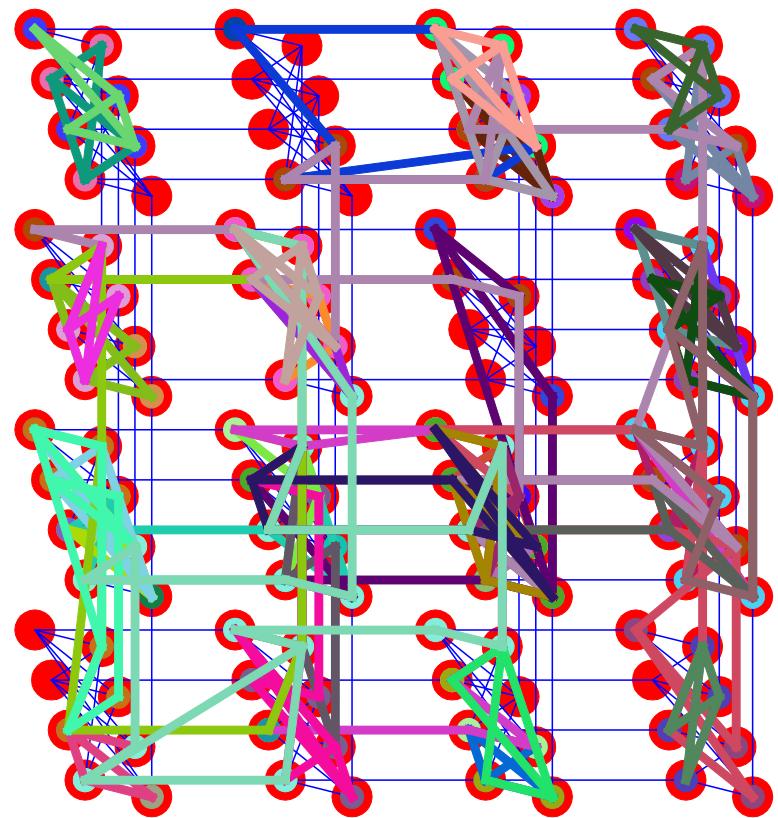
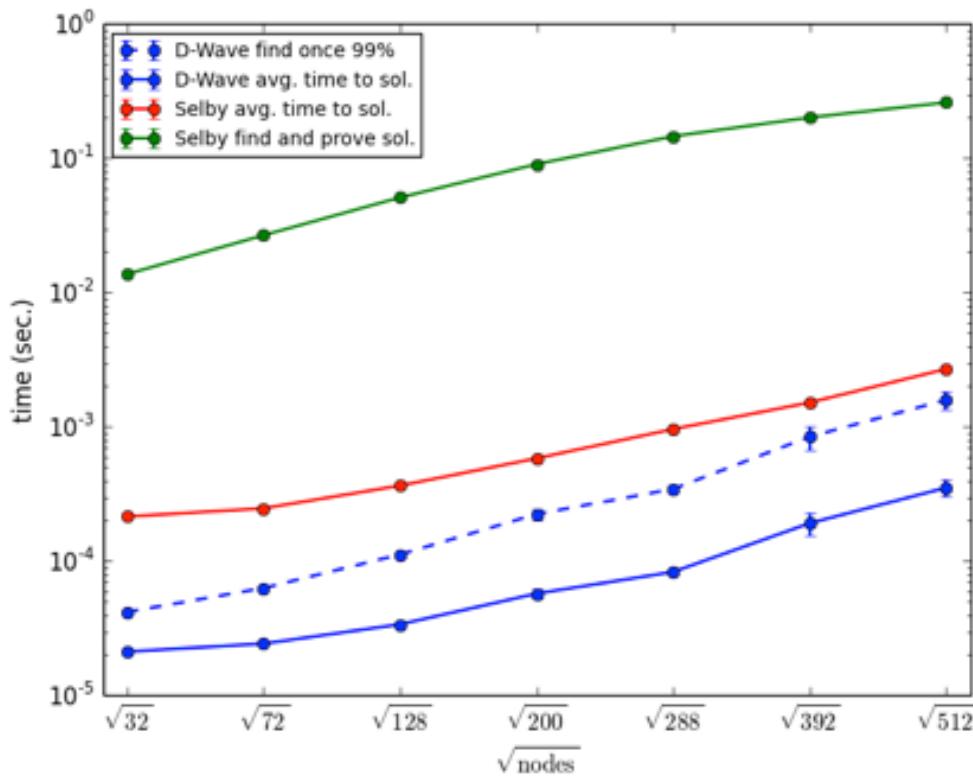
Example: easy problems



- Independent Set (IS) [left]: fundamental comb. opt. problem
- Affinity Independent Set (AIS) [right]: new social network analysis variant
- IS solvable in polynomial time on Chimera, yet appears hard for D-Wave
- AIS likely solvable in polynomial time too, yet appears easy for D-Wave
- AIS is just IS with -1,+1 linear-term weights rather than 0,1 as for IS
- Illustrates difficulties in empirical benchmarking vs theoretical analysis & impact of problem formulation



Example: open problem



- “Frustrated loop” instances by Itay Hen et al. [AQC Workshop, 2014]
- Potential example of a “win” for D-Wave
- However, we show solution *value* can be computed in polynomial time

Configuration matters

- Random instances on D-Wave hardware
 - $\{-1,1\}$ coupler values
 - D-Wave Two finds optimal in 0.5 sec, while classical algorithms scale poorly
[McGeoch and Wang, [Conf. Computing Frontiers 2013: 23](#)]
 - Claimed 3600x speedup
 - We observe classical Integer Program solvers match performance with appropriate model
[also Dash, [arXiv:1306.1202v2](#) (2013)]
 - Speedup vanishes with proper configuration/usage
 - QUBO vs Ising
 - Random instance with linear term appear easy for former but hard for latter!

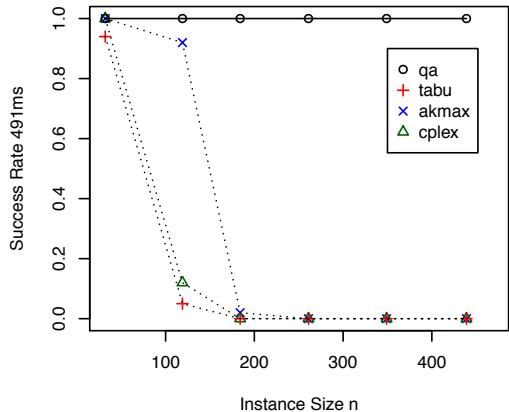


Figure 1: Success rates: proportion of best solutions found in 491ms CPU time (tabu, amax, cplex software) and exclusive access time (QA hardware).

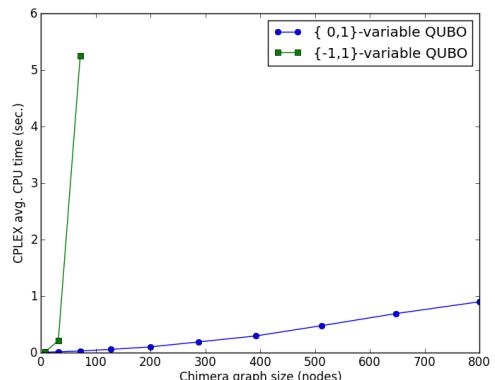
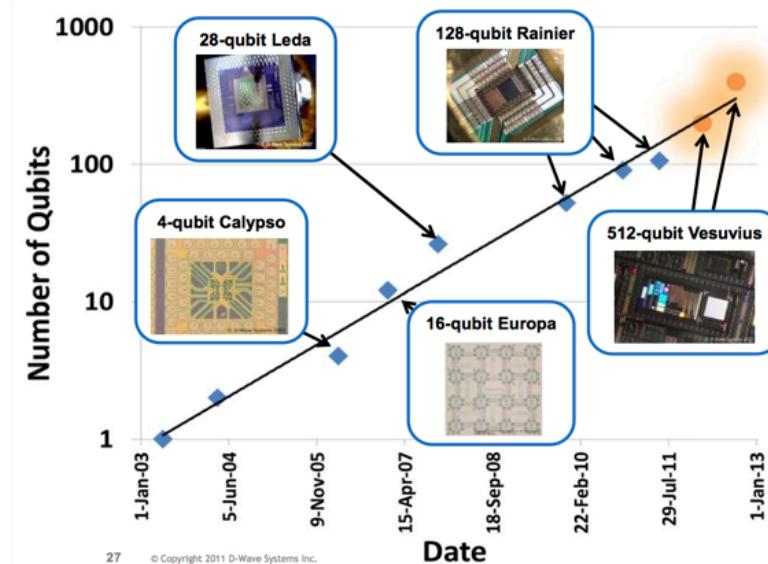
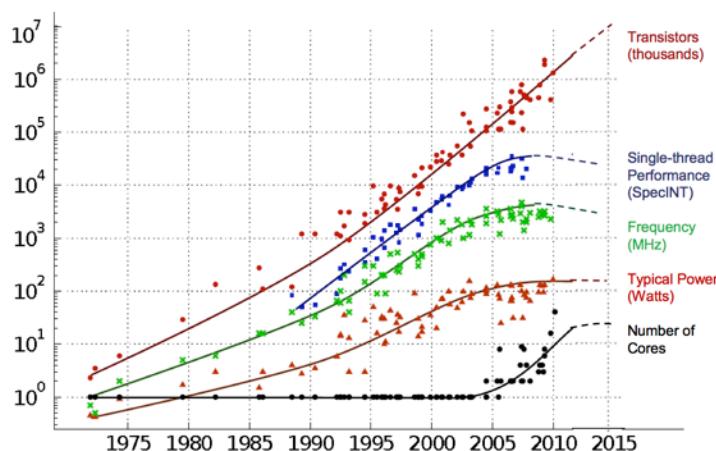
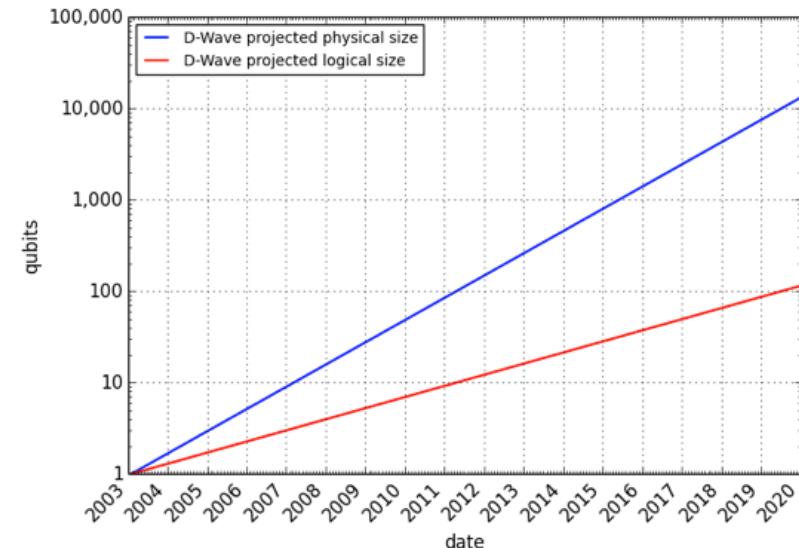


Figure 2: We observed CPLEX performance on random $\{0,1\}$ -variable QUBO can match D-Wave performance as reported above by McGeoch and Wang.

D-Wave system scaling



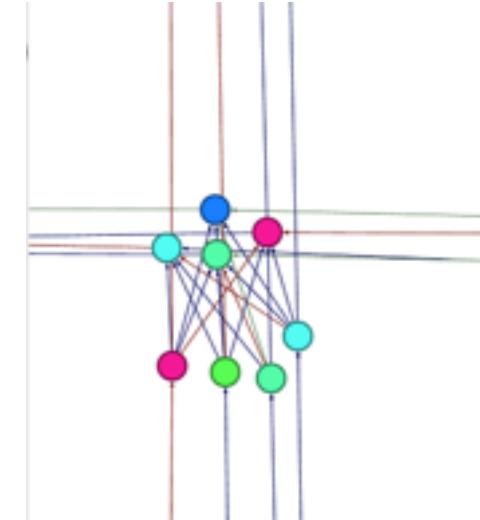
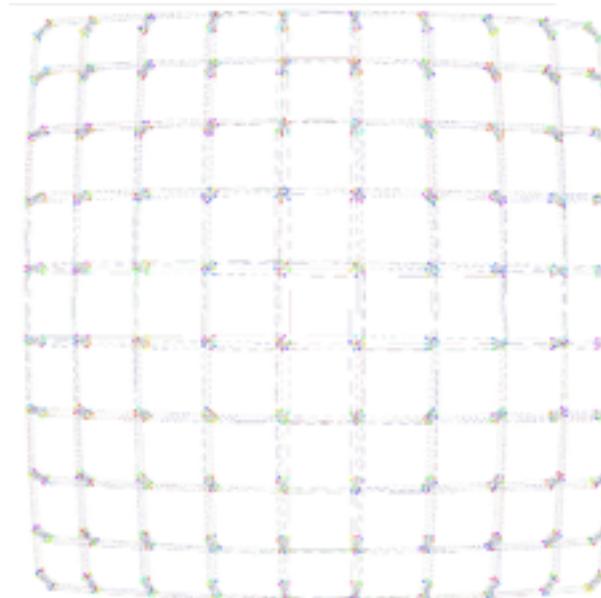
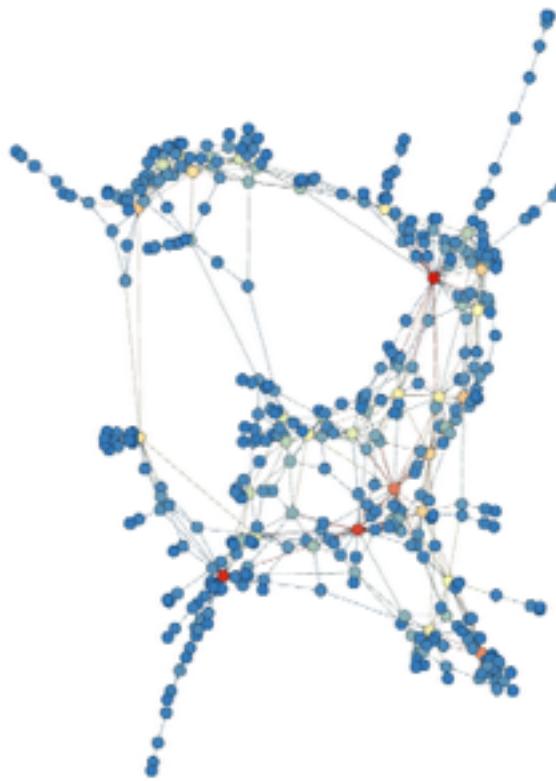
27 © Copyright 2011 D-Wave Systems Inc.



Original data collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond and C. Batten
Dotted line extrapolations by C. Moore

- D-Wave scaling projections are promising (relative to conventional computing)
- However, Chimera architecture is significant bottleneck [top]
- Algorithmic tools for representing real-world problems on emerging quantum architectures are critical!
- We initiate such an endeavor

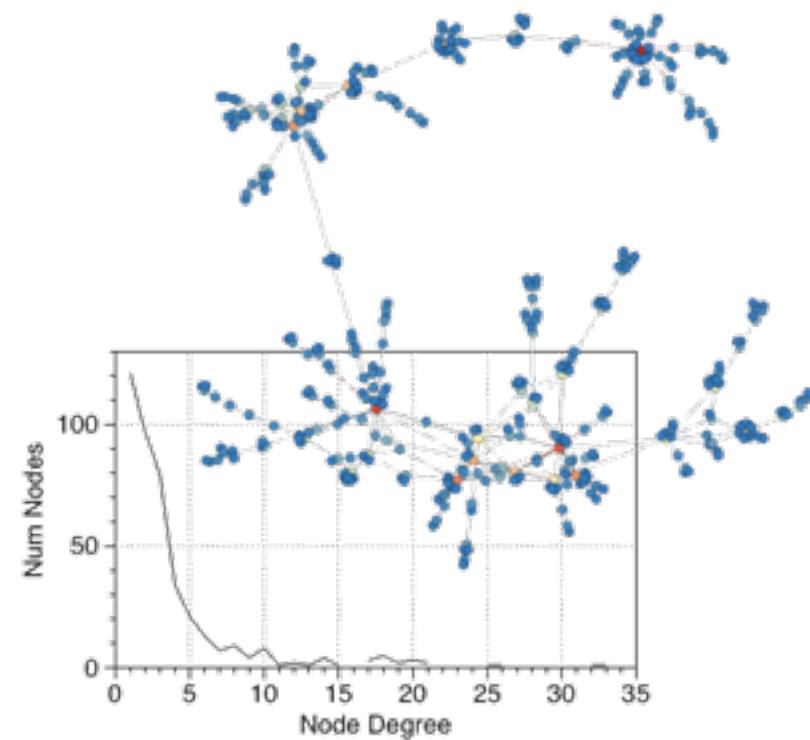
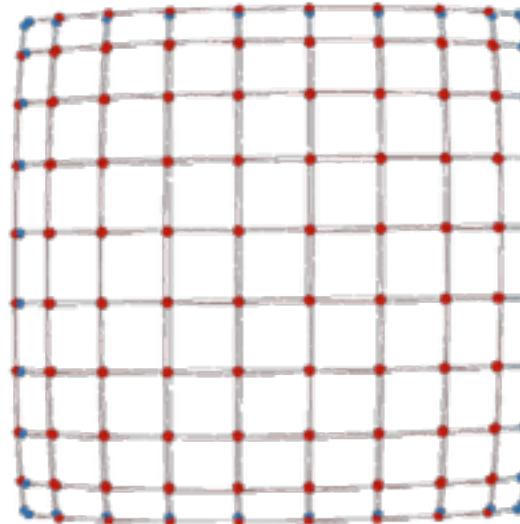
Complex networks on the Chimera graph



- New approach to circumvent embedding
- Generate complex network simultaneously while embedding it
- Efficiency for 512-node Chimera around 50% vs 6.25% worst case

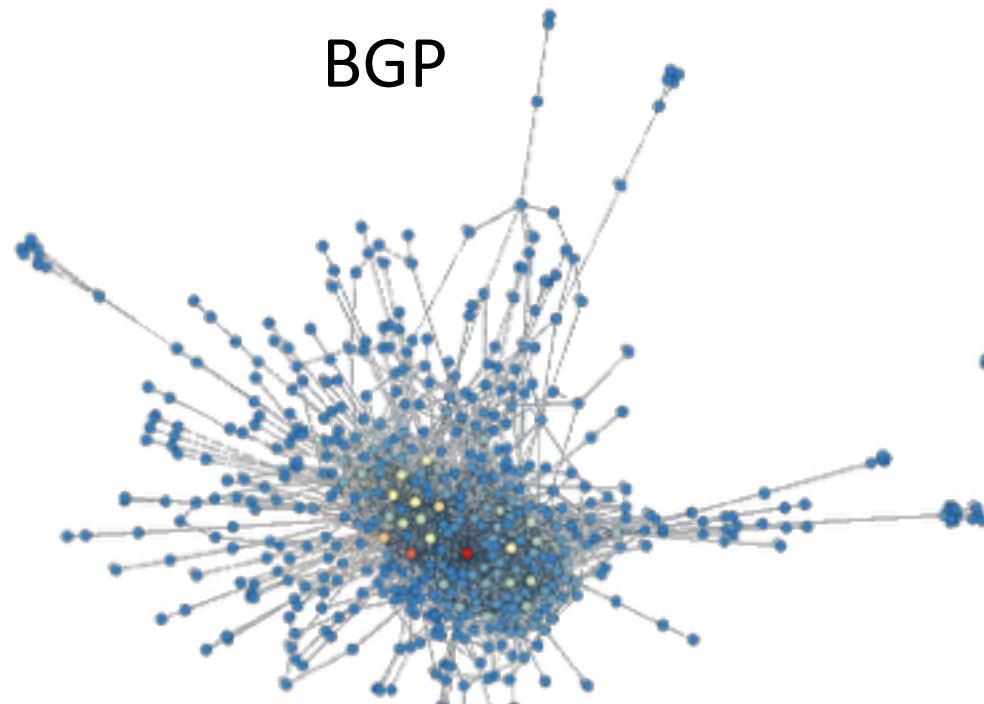
Generating complex networks on a Chimera graph

- Mapping arbitrary graph to Chimera is hard
- Instead, alter Chimera graph to have “real-world” properties
 - Merge nodes to increase node degree
 - Remove edges between nodes to selectively decrease degree

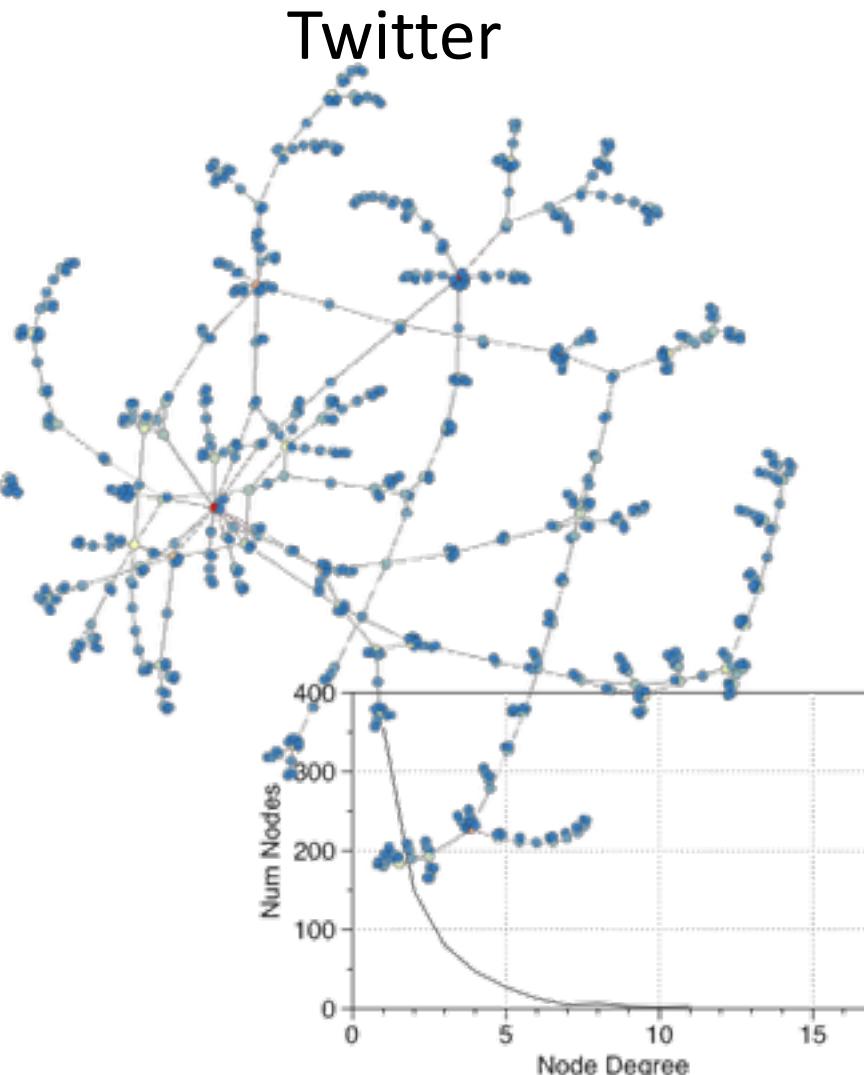


Real-world complex networks

BGP

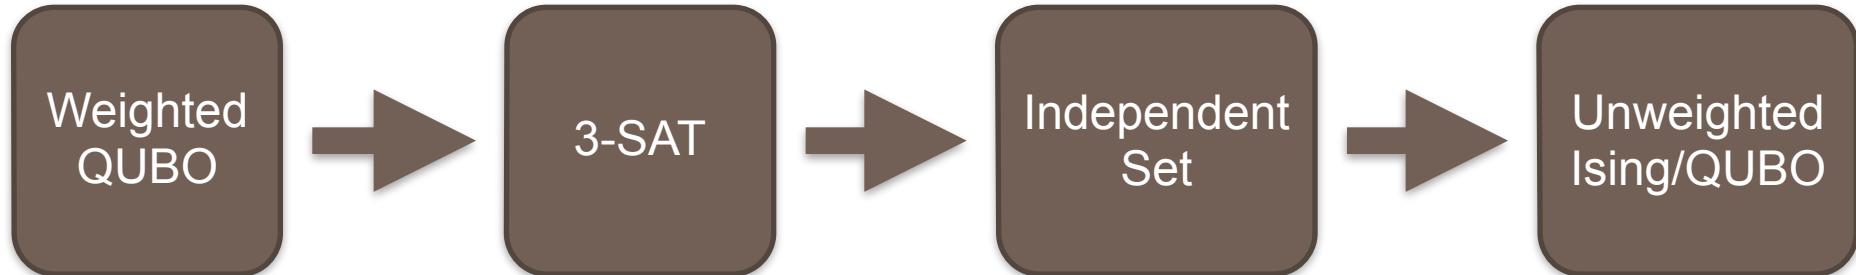
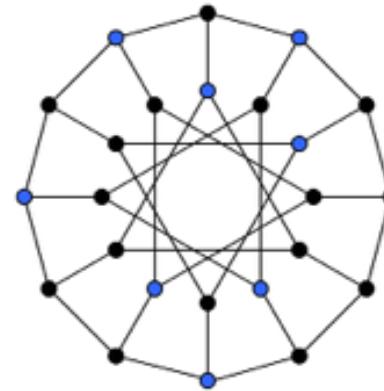


Twitter



From weighted to unweighted QUBO

$$\min_{x_i \in \{0,1\}} \sum_{ij} A_{ij} x_i x_j + \sum_i c_i x_i$$



$$\Phi_1(\vec{x}, \vec{y}) = \bigwedge_{ij \in E} (\overline{y_{ij}} \vee x_i \vee x_j) \wedge (\overline{y_{ij}} \vee \overline{x_i} \vee \overline{x_j}) \wedge \\ (y_{ij} \vee x_i \vee \overline{x_j}) \wedge (y_{ij} \vee \overline{x_i} \vee x_j)$$

$$\min_{x_i \in \{-1,1\}} \sum_{ij \in E'} x_i x_j - \sum_{i \in V'} x_i$$

This reduction converts a weighted instance into an unweighted but larger instance, enabling a trade off between coupler precision and qubits.

Contributions

- Brought real-world problems and instances to bear
- Developed more rigorous benchmarking methodology
 - Exposed importance of success criterion and algorithm selection and configuration
 - Drawing from theoretical CS techniques and results
 - Demonstrated formal hardness of widely used instances
- Techniques for circumventing D-Wave architectural limitations
 - Novel technique for sidestepping embedding issues by generating social network minors within Chimera graph
 - Trading off coupler precision and qubits via NP reductions
- New QUBO formulations for complex network analysis problems

Concluding thoughts on D-Wave



- Even if D-Wave's machines are truly quantum computers, the user is shielded from the quantumness
- Convenient to think of D-Wave as a QUBO-solving black box that implements a particular meta-heuristic (e.g. akin to simulated annealing)
- With the above perspective: benchmark and analyze D-Wave performance as we do with other meta-heuristics
- Fair point of comparison would be against, e.g., best classical (approximate/randomized) QUBO solvers
- Limited Chimera connectivity and coupler precision are most significant barriers to solving interesting and real-world instances
- “Arms race” between classical and D-Wave/quantum is healthy

Supplementary slides

Complexity of QUBO on Chimera

- (Decision version) NP-complete even with no linear term and $\{-1, 0, 1\}$ weights [Barahona, 1982]
- We show NP-complete with no linear term and $\{-1, 1\}$ weights
 - Instances used in D-Wave benchmarking studies
- Tree-width (path-width) is $\Theta(\sqrt{n})$, yielding $O(2^{\sqrt{n}})$ algorithm
 - “Subexponential” exact algorithm even though NP-hard
- Approximation complexity?
 - Polynomial-time approximation scheme (PTAS)
[Saket, 2013, [arXiv:1306.6943](https://arxiv.org/abs/1306.6943)]
 - PTAS’s are rarely efficient; theory vs practice?
 - Efficient approx algorithm for say, getting within 90%?

Limits of reducing to Chimera

- Can we do better than a quadratic blowup in qubits?
- Probably not, due to Exponential Time Hypothesis
 - Problems like Max-Cut on general graphs are conjectured to require $O(2^n)$ time
 - But we have a $O(2^{\sqrt{n}})$ time algorithm for Chimera Ising
 - So in some sense quadratic factor is artifact of Chimera
- Weights make this worse: Choi embedding assumes (linearly) large weights
- Reduction better than $O(n^2)$ for Max-Cut on bounded-degree graphs would improve best-known classical algorithm
 - Applies to **any** reduction, not just minor embeddings

Max-Cut as model problem

- NP-complete on bipartite graphs with weights {-1,1}
 - Replace each edge by a path of +1 and -1 edge
- Max-Cut essentially equivalent to Ising problem
 - Can use an apex vertex to model linear term
- We give reduction from weighted QUBO to unweighted QUBO
 - Unroll and optimize existing chain of reductions
 - Weights are a significant barrier in D-Wave benchmarking
- Our reduction only incurs linear blowup on bounded-degree graphs
 - However, does not preserve Chimera structure

Weighted QUBO as 3-SAT

x_1, \dots, x_n



$$\Phi_1(\vec{x}, \vec{y}) = \bigwedge_{ij \in E} (\overline{y_{ij}} \vee x_i \vee x_j) \wedge (\overline{y_{ij}} \vee \overline{x_i} \vee \overline{x_j}) \wedge \\ (y_{ij} \vee x_i \vee \overline{x_j}) \wedge (y_{ij} \vee \overline{x_i} \vee x_j)$$

$$y_{ij} = x_i \oplus x_j$$

k, w_1, \dots, w_m



y_1, \dots, y_m



Adder circuit: $\sum_i w_i y_i \geq k$?

$\Phi_2(\vec{y}, \vec{w}, k)$



$\{0, 1\}$

$\Phi = \Phi_1 \wedge \Phi_2$

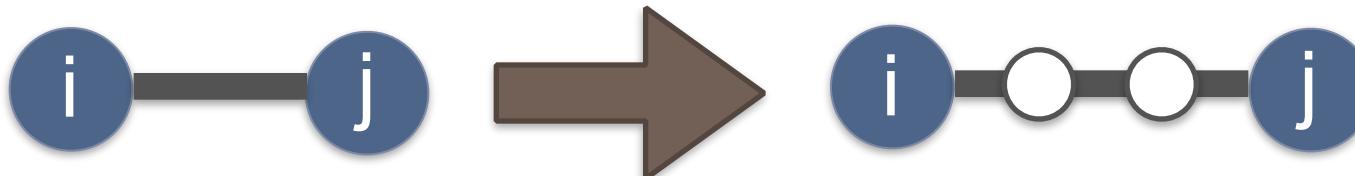
3-SAT as unweighted Ising

- Standard reduction from 3-SAT to Independent Set
- Usual Ising formulation of Independent Set:

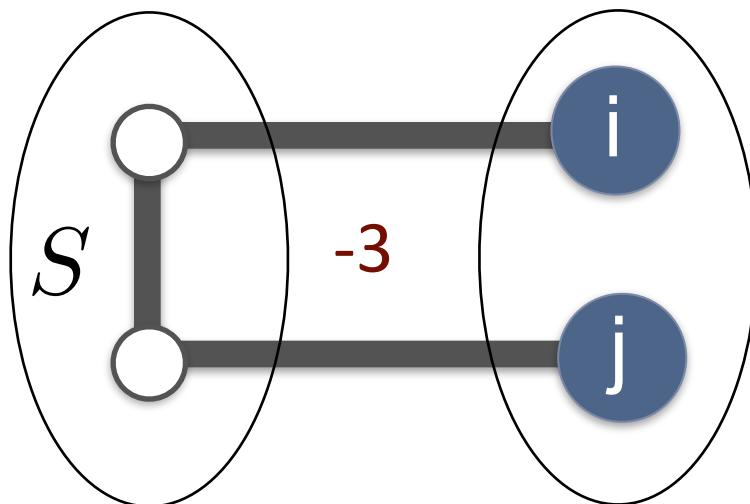
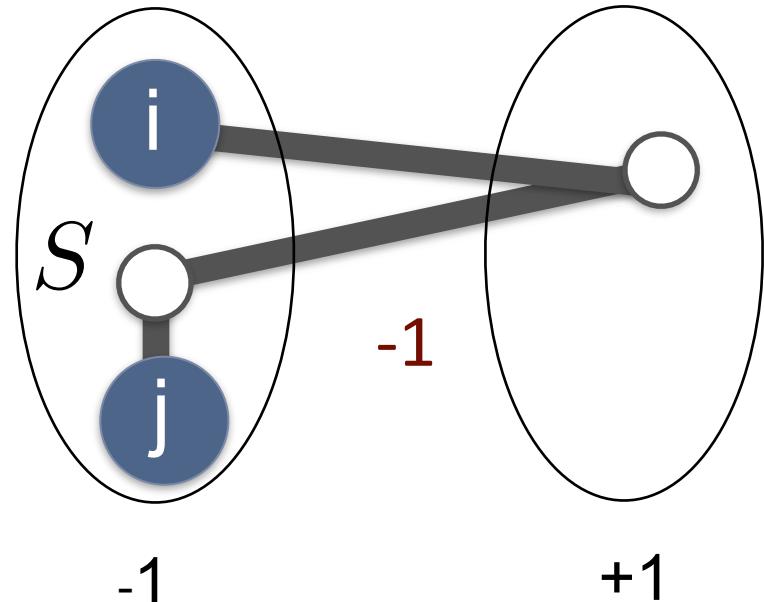
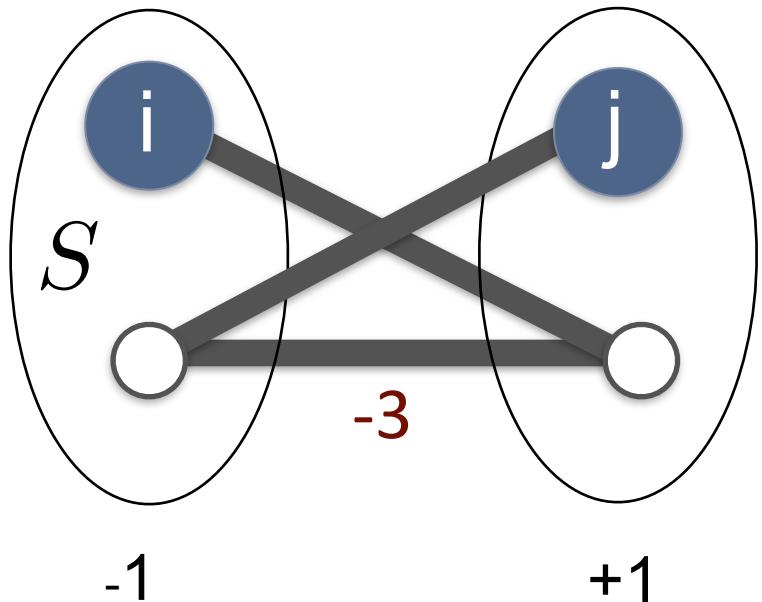
$$\min_{x_i \in \{-1,1\}} \sum_{ij \in E} x_i x_j - \sum_{i \in V} d_i x_i$$

- New formulation on graph where each edge replaced by 3-path circumvents linearly large weights above:

$$\min_{x_i \in \{-1,1\}} \sum_{ij \in E'} x_i x_j - \sum_{i \in V'} x_i$$



Independent set as Ising



WLOG: -1 side is an independent set

Objective value:

$$-|S| + (|V| - |S|) - 3|E| = |V| - 2|S| - 3|E|$$