Adiabatic Quantum Annealing Update

Bob Lucas
USC – Lockheed Martin Quantum Computing Center
July 8, 2014

USC Colleagues

Information Sciences Institute

Itay Hen

Rosa Difilice

Federico Spedalieri

Greg Ver Steeg

Daniel Lidar with Graduate Students and Post Docs

Prof. Daniel Lidar

Zhihui Wang

Tameem Albash

Kristen Pudenz

Joshua Job

Anurag Mishra

Broader Collaborators

Lockheed Martin

Greg Tallant, Peter Stanfill, and Steve Adachi

ETH

Matthias Troyer and Troels Ronnow

Google

Sergio Boixo

UCL

Paul Warburton and Walter Vinci

Outline

D-Wave Machine

Research Results to Date

Applications

Summary

.

The End of Dennard Scaling

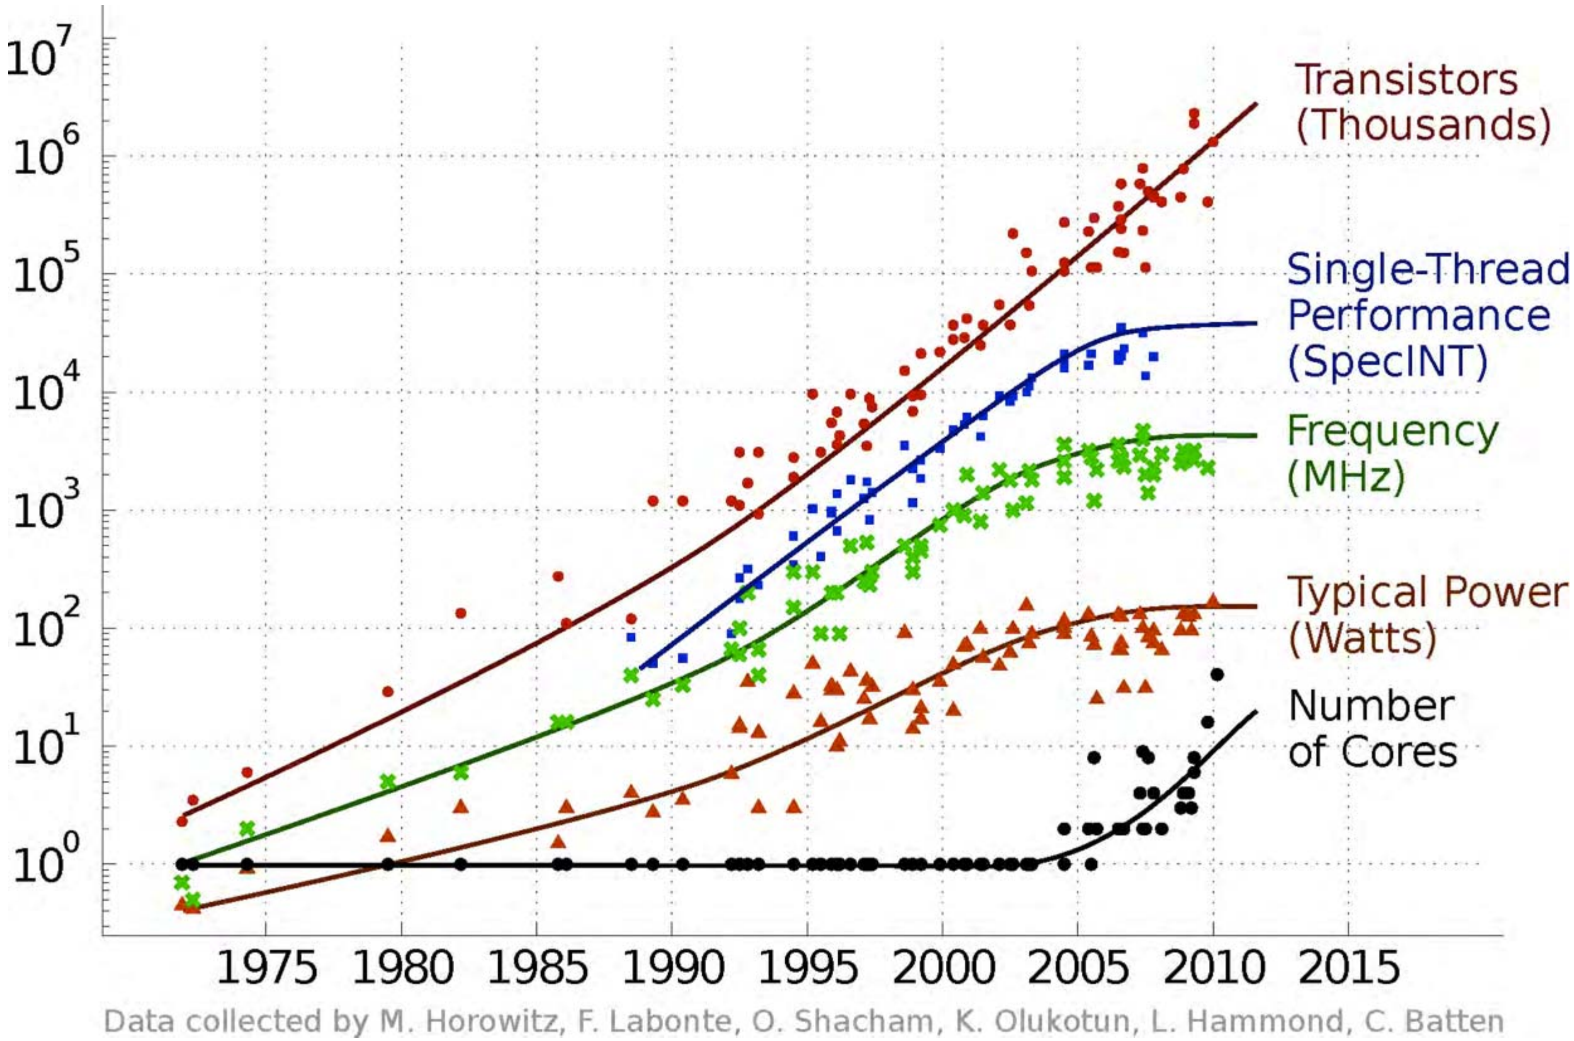

Need More Capability?

Massive Scaling
Tianhe-2 (3M cores)

Exploit a New Phenomenon
D-Wave Quantum Annealer

Application Specific Systems
D.E. Shaw Research Anton

Adiabatic Quantum Annealing

Problem: find the ground state of

$$H_{\text{Ising}} = \sum_j h_j \sigma_j^z + \sum_{(i,j)\in E} J_{ij} \sigma_i^z \sigma_j^z$$

Shown by Barahona (1982) to be NP-hard in 2D, $J_{ij} = \pm,\ h_j \neq 0.$

Use adiabatic interpolation from transverse field (Farhi et al., 2000)

$$H(t) = A(t)\sum_j \sigma_j^x + B(t)H_{\text{Ising}}$$

$t \in [0, t_f]$

Program $\{h_i\},\{J_{ij}\}$

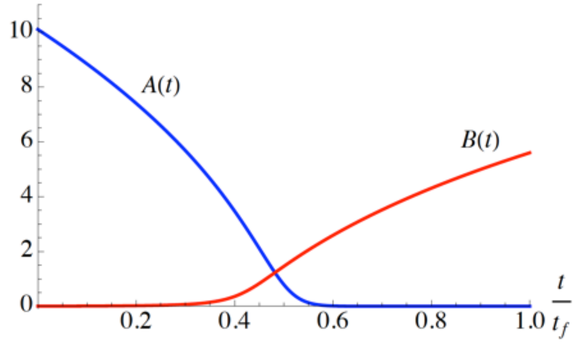

Graph Embedding implemented on DW-1 via Chimera graph retains NP-hardness (V. Choi, 2010)

D-Wave System Collage

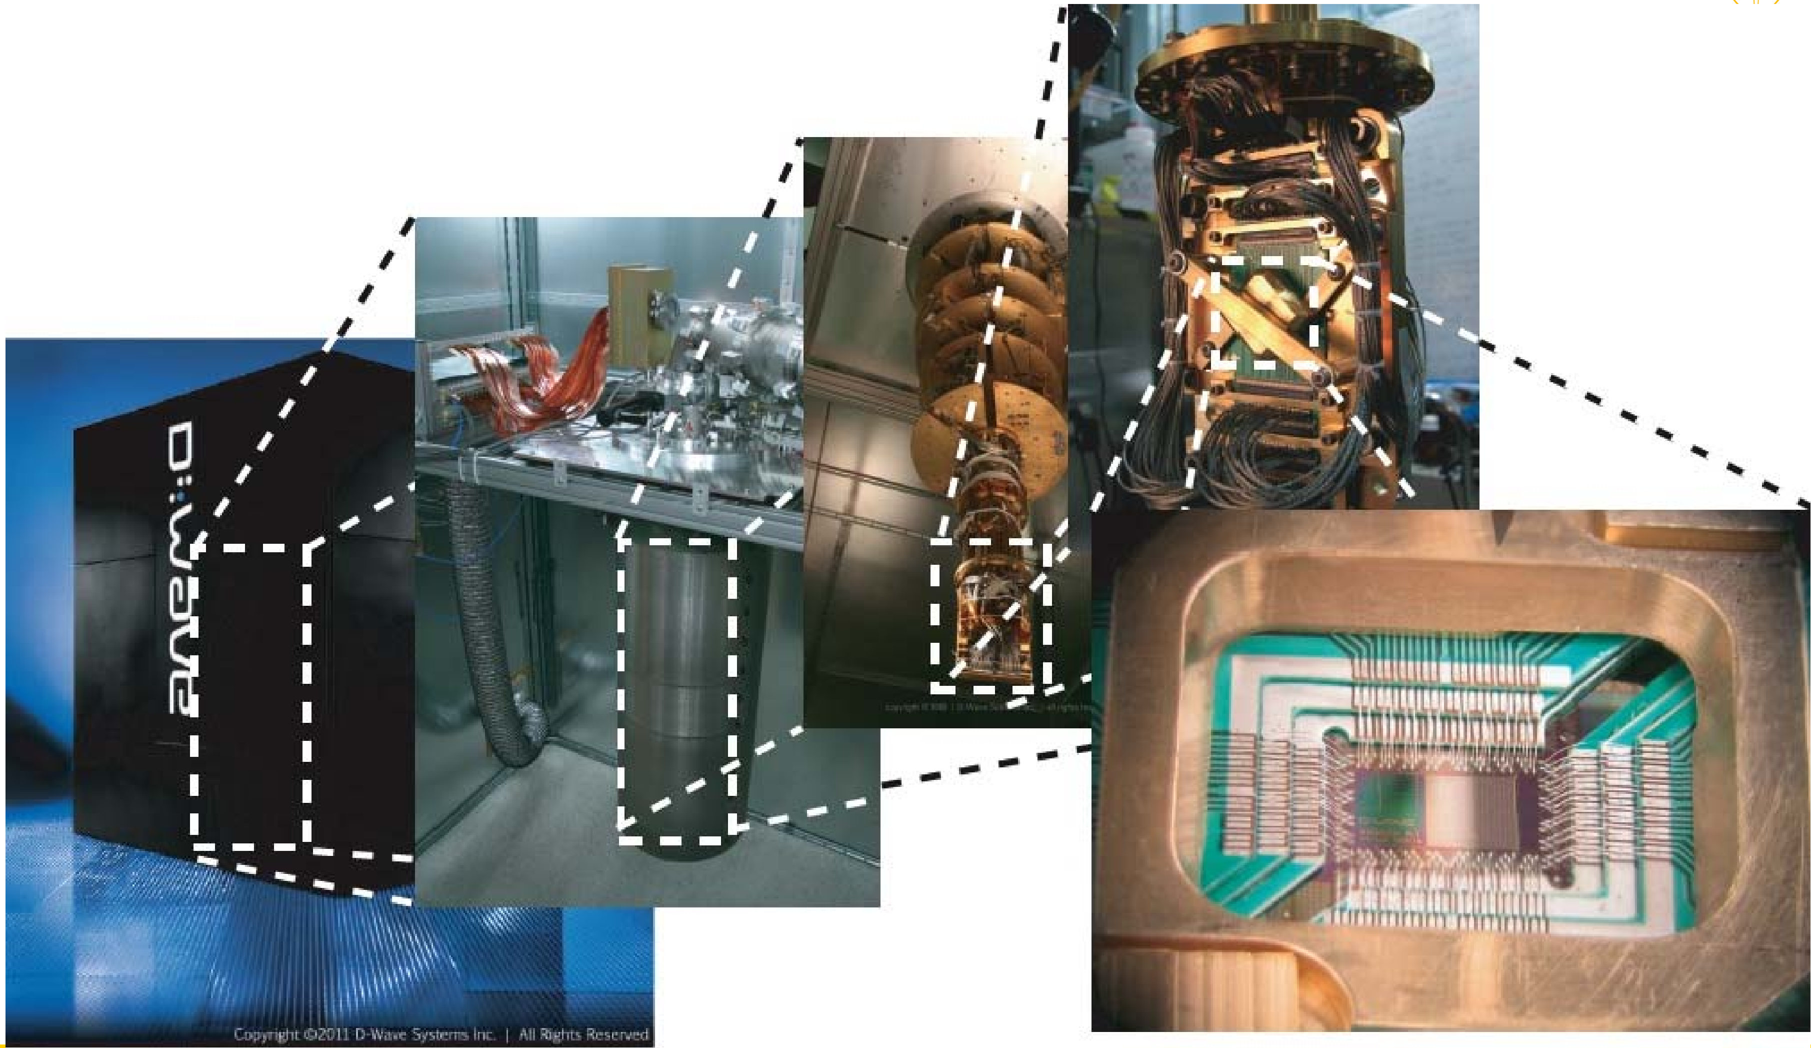

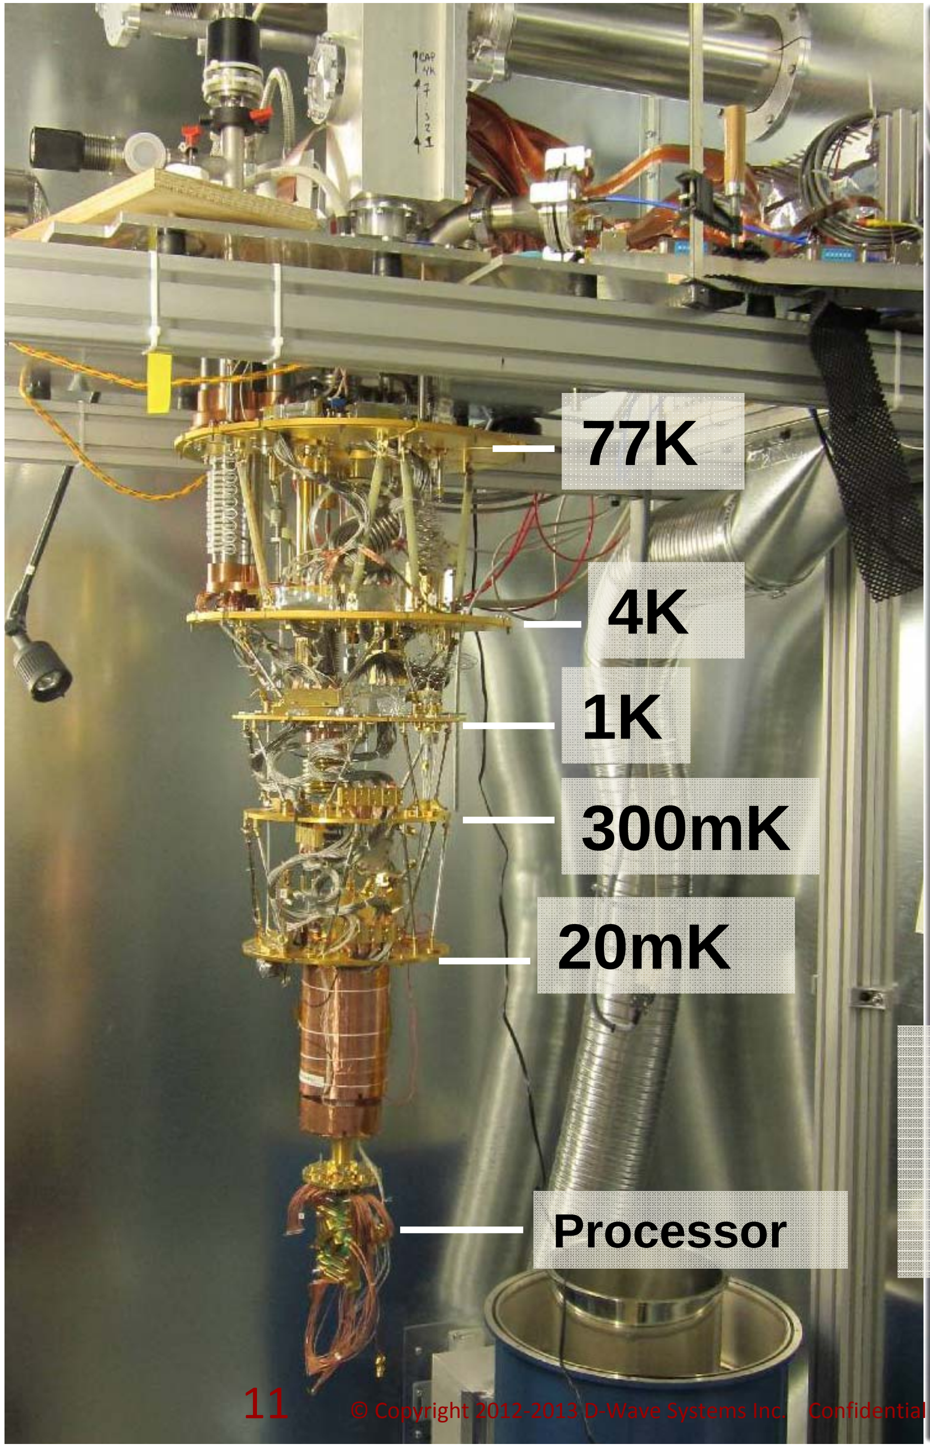
77K
4K
1K
300mK
20mK
Processor

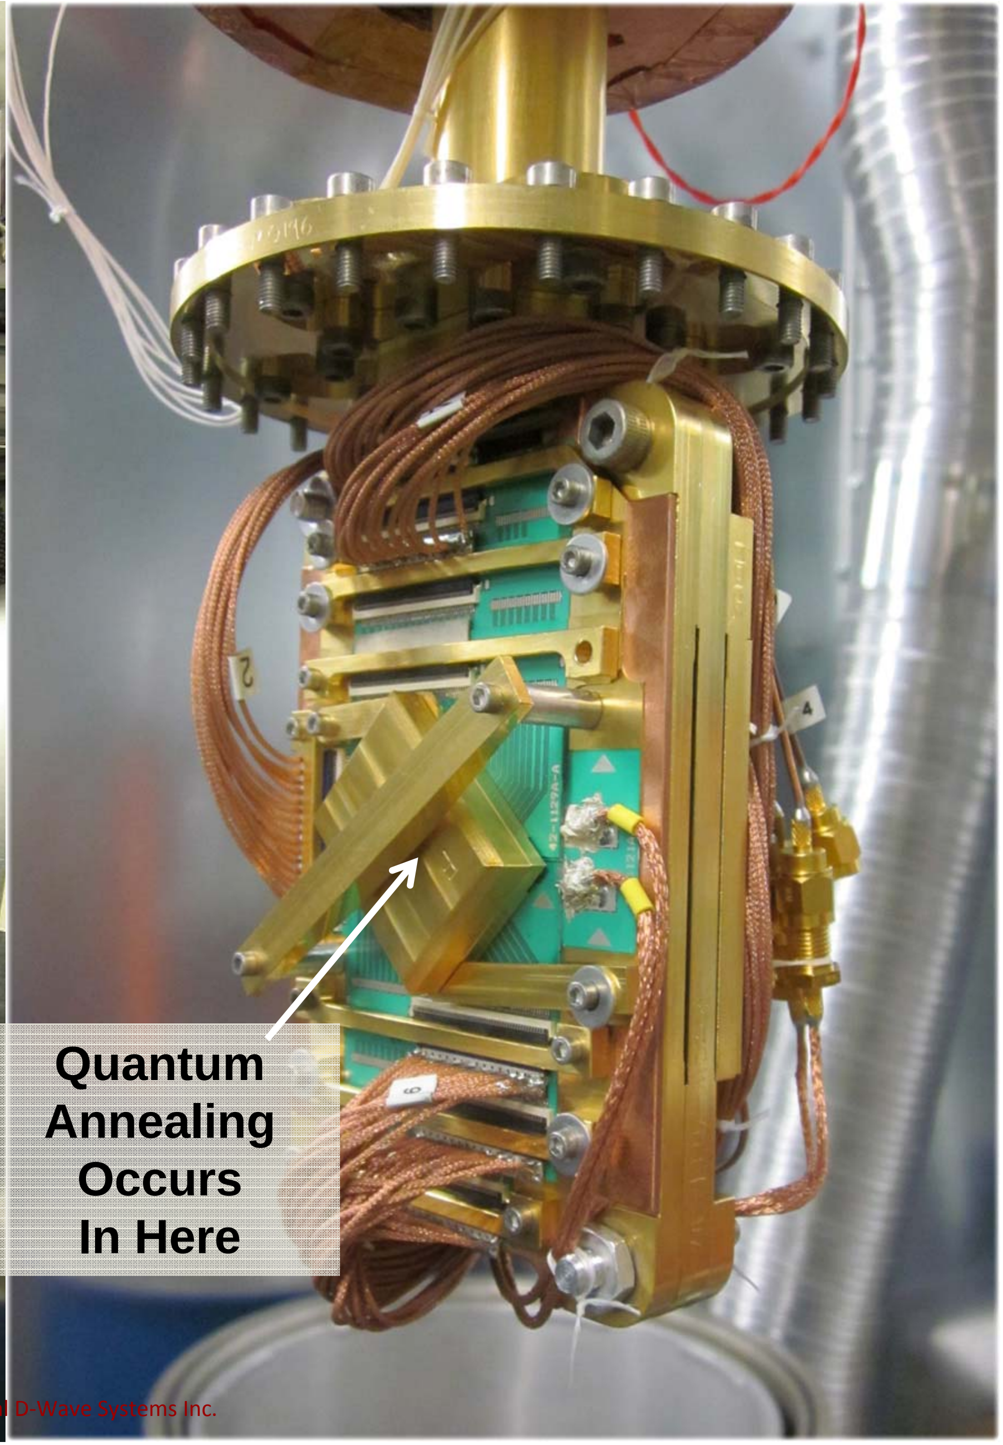
Quantum Annealing Occurs In Here

Eight Qubit Unit Cell

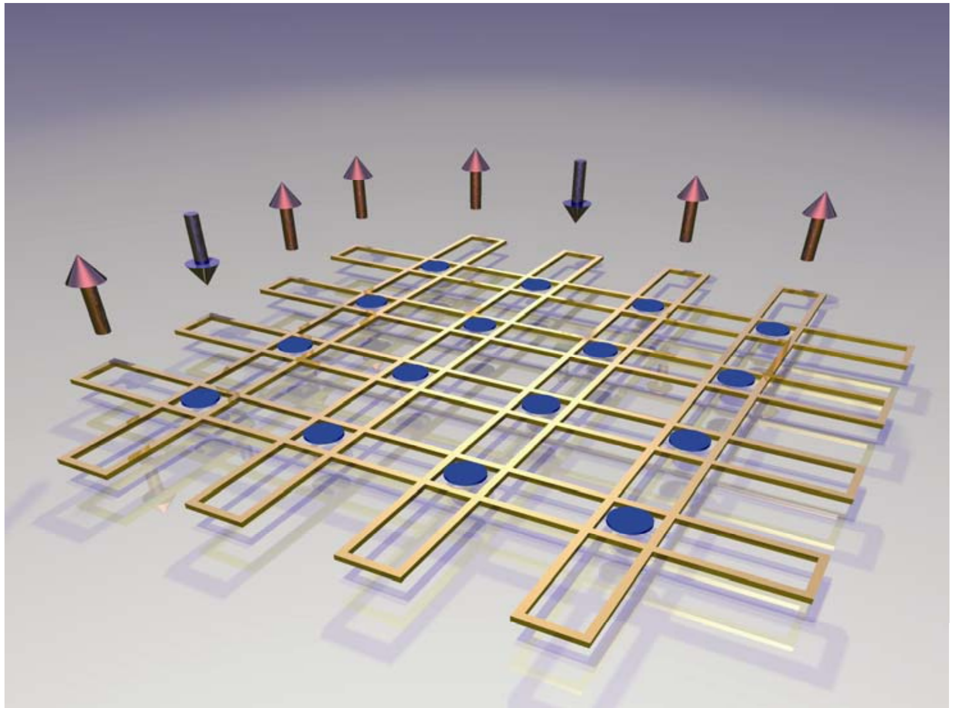

Images courtesy D-Wave

Chimera Graph Topology

The topology of the D-Wave Two at the USC – Lockheed Martin Quantum Computing Center.

503 of 512 qubits calibrated and mapped.

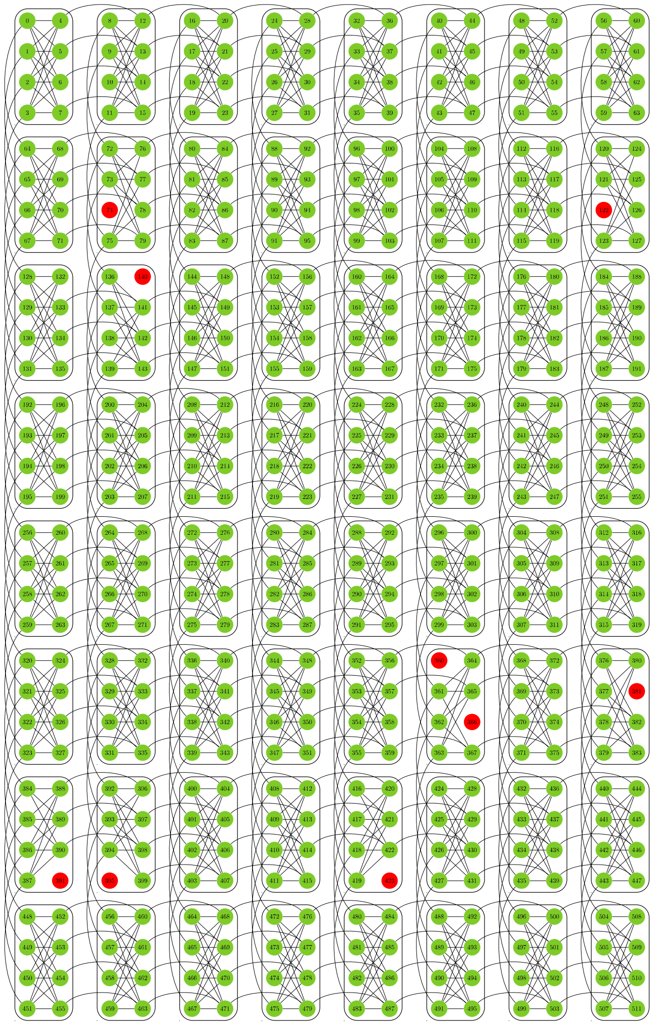

D-Wave's Version of "Moore's Law"

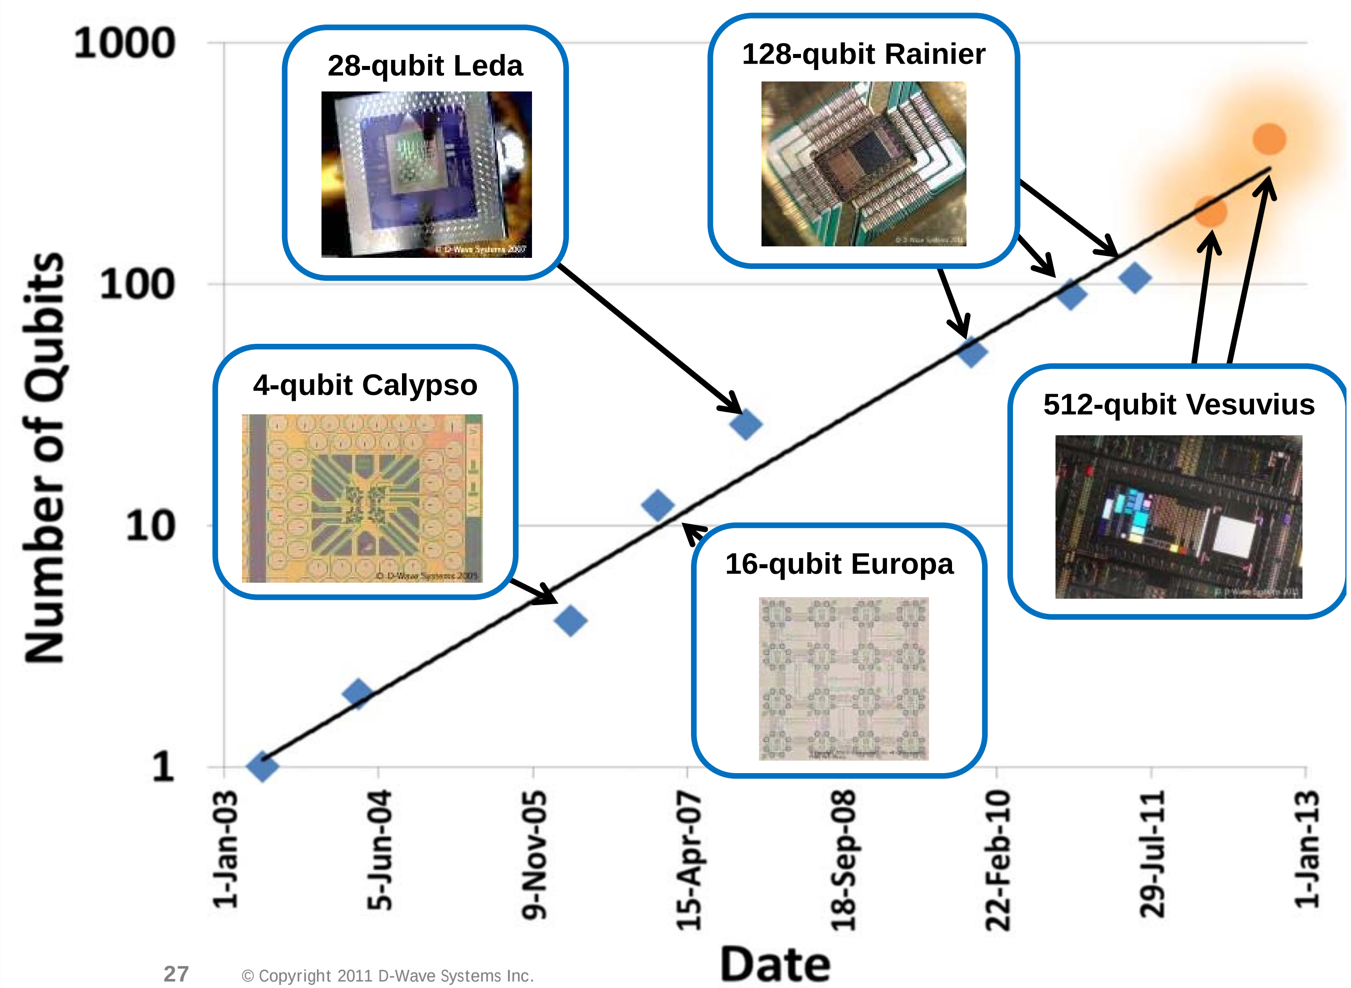

Outline

D-Wave Machine

Research Results to Date

Applications

Summary

.

Quantum Signature: Degenerate Ising Hamiltonian

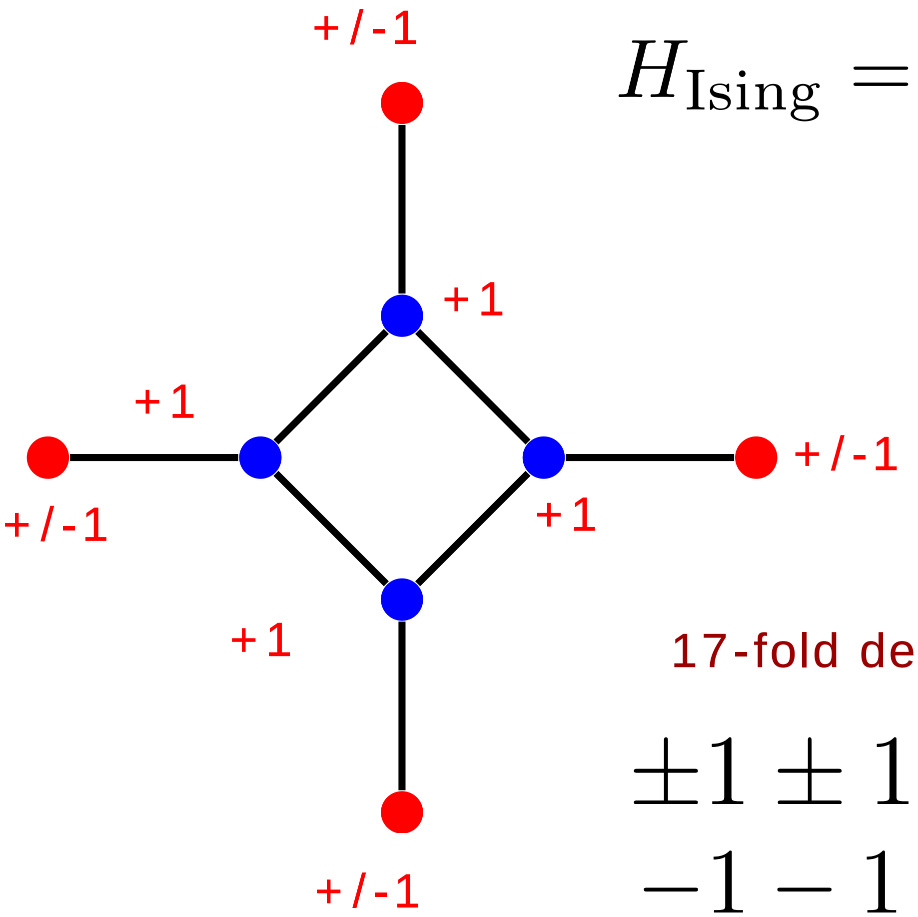

$$H_{\text{Ising}} = \sum_j h_j \sigma_j^z + \sum_{(j,k)\in E} J_{jk} \sigma_j^z \sigma_k^z$$

hj = -1, hj = 1, Jik = -1

17-fold degenerate ground space:

$$\begin{array}{cccccccc} \pm 1 & \pm 1 & \pm 1 & \pm 1 & 1 & 1 & 1 & 1 \\ -1 & -1 & -1 & -1 & -1 & -1 & -1 & -1 \end{array}$$

Simulated Annealing At Several Speeds

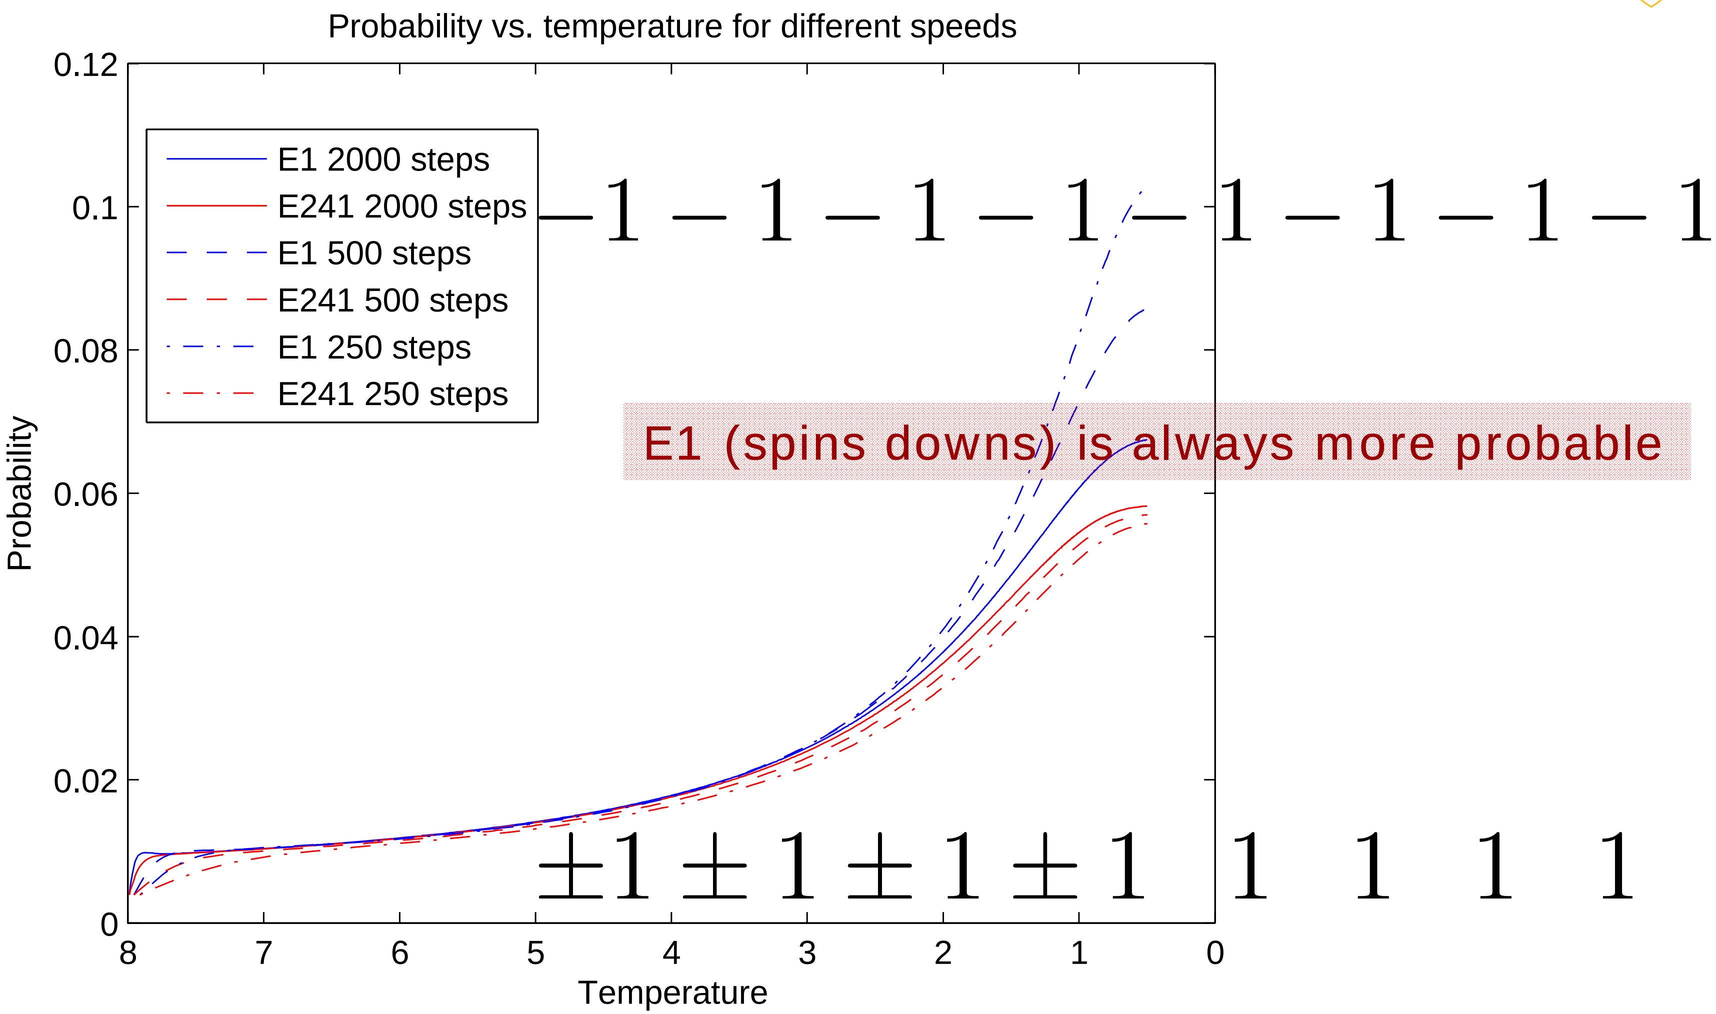

E1 (spins downs) is always more probable

Simulated Quantum Annealing

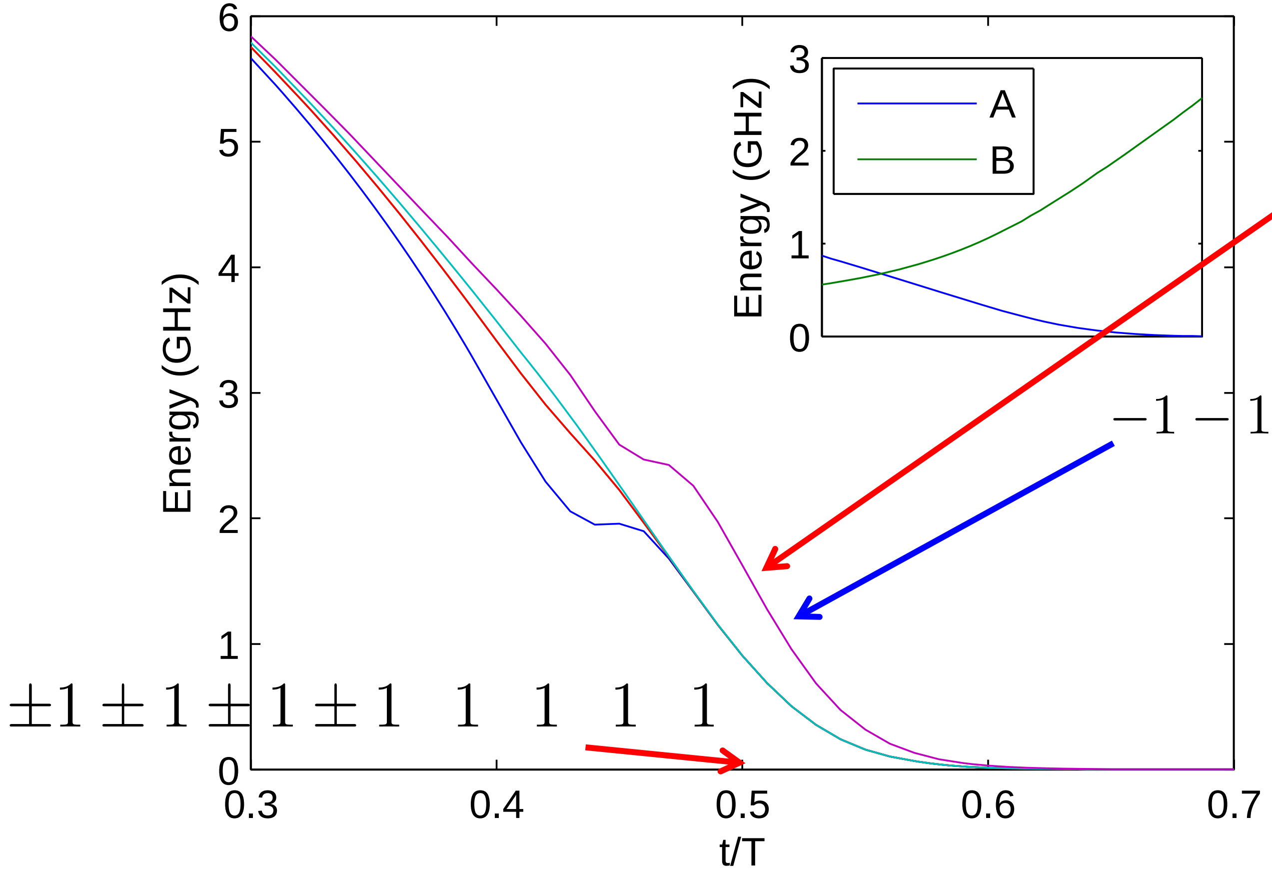

Gap 1.5 GHz

(Temp: 0.35 GHz)

$-1\ -1\ -1\ -1\ -1\ -1\ -1\ -1$

D-Wave One Experiments

144 embeddings

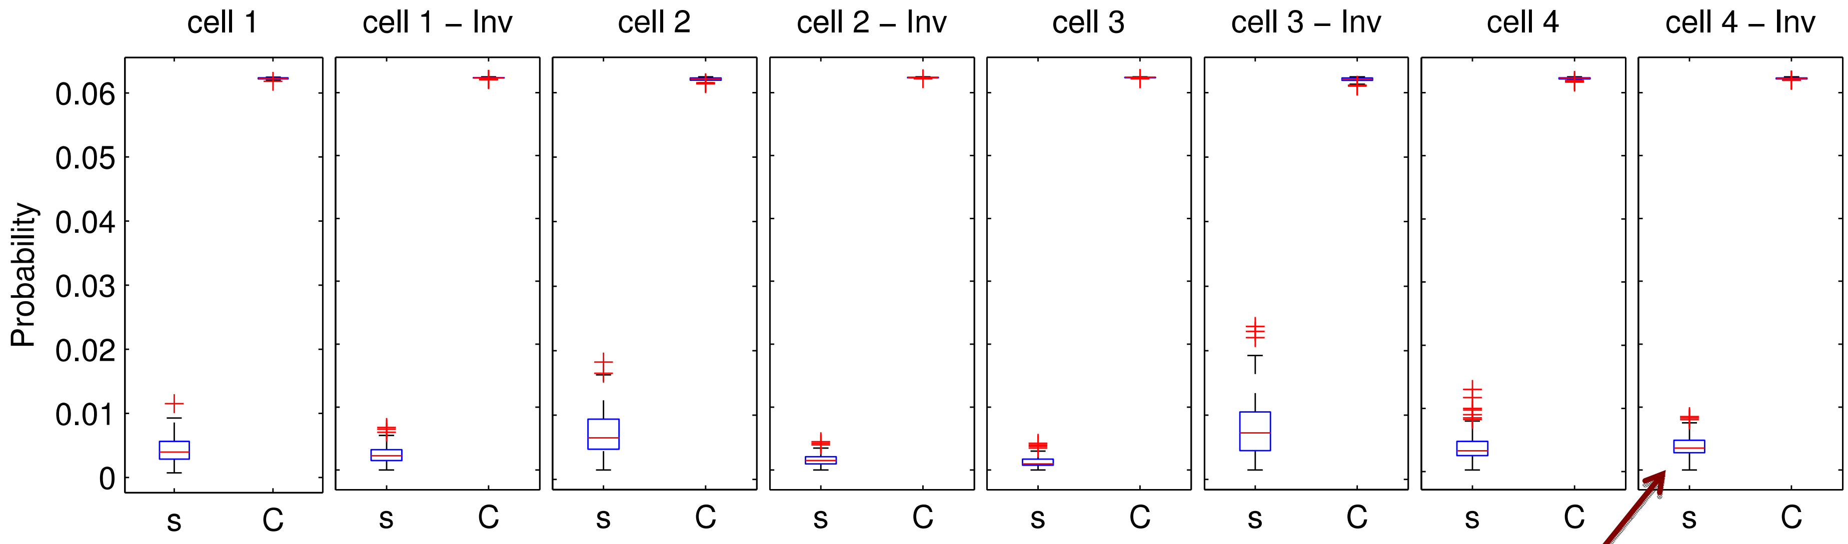

Quantum Signature: the isolated state is suppressed

Reference: S. Boixo,, T. Albash, F.M. Spedalieri, N. Chancello, D.A. Lidar, *Nature Comm*. **4**, 2067 (2013).

Scalable Quantum Signature

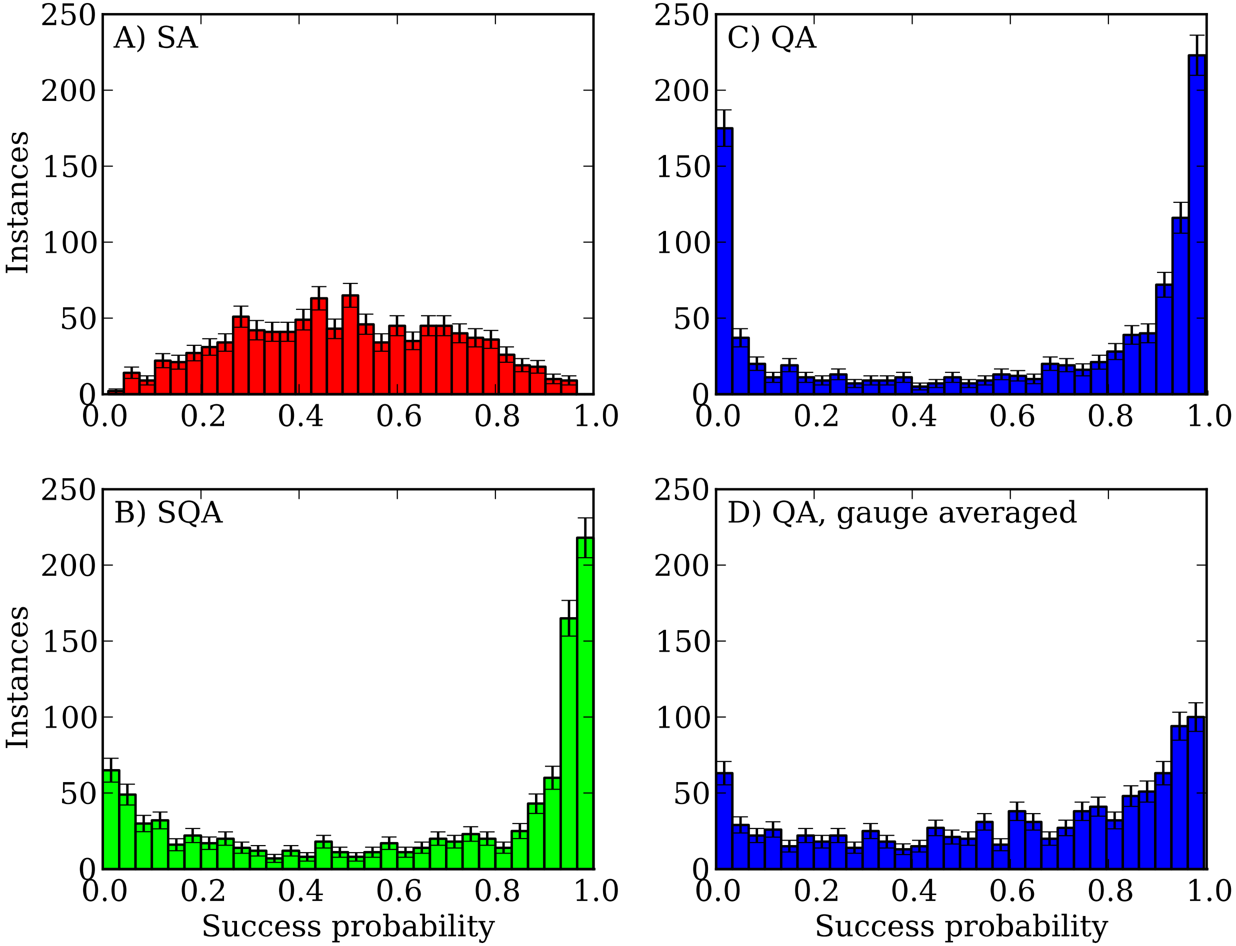

S. Boixo, T. F. Rønnow, S. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, M. Troyer, *Nature Physics* **10**, 218–224 (2014)

Witness for Entanglement

Collaboration with D-Wave

Only they can take the tunneling microscopy measurements needed
Allows us to infer the value of an entanglement witness

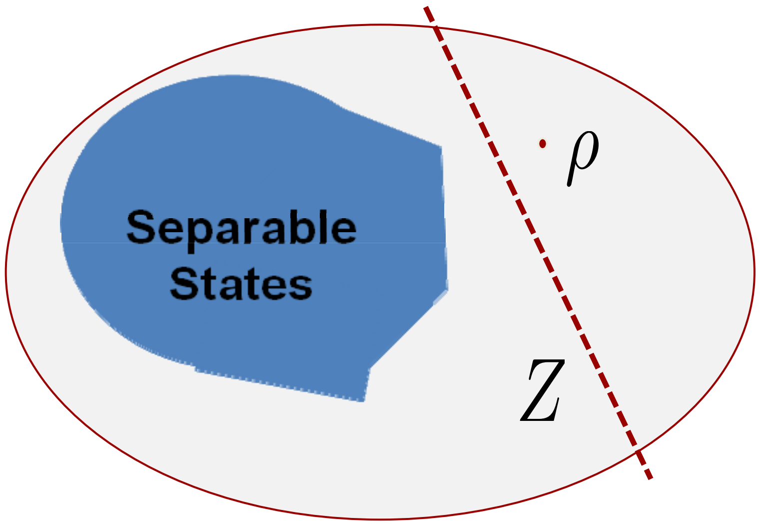

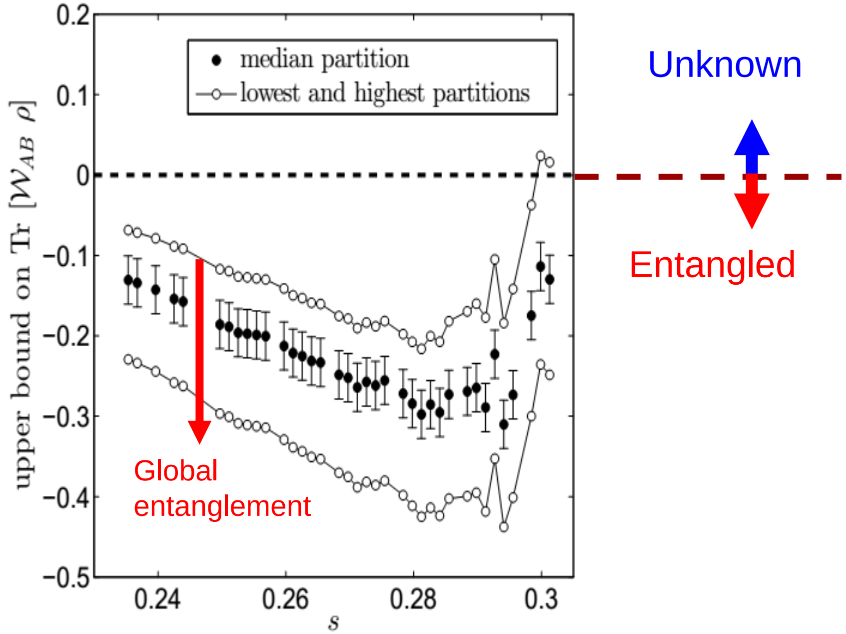

Lanting, et.al, PRX 4, 021001 2014

Performance of Classical Exact Solvers

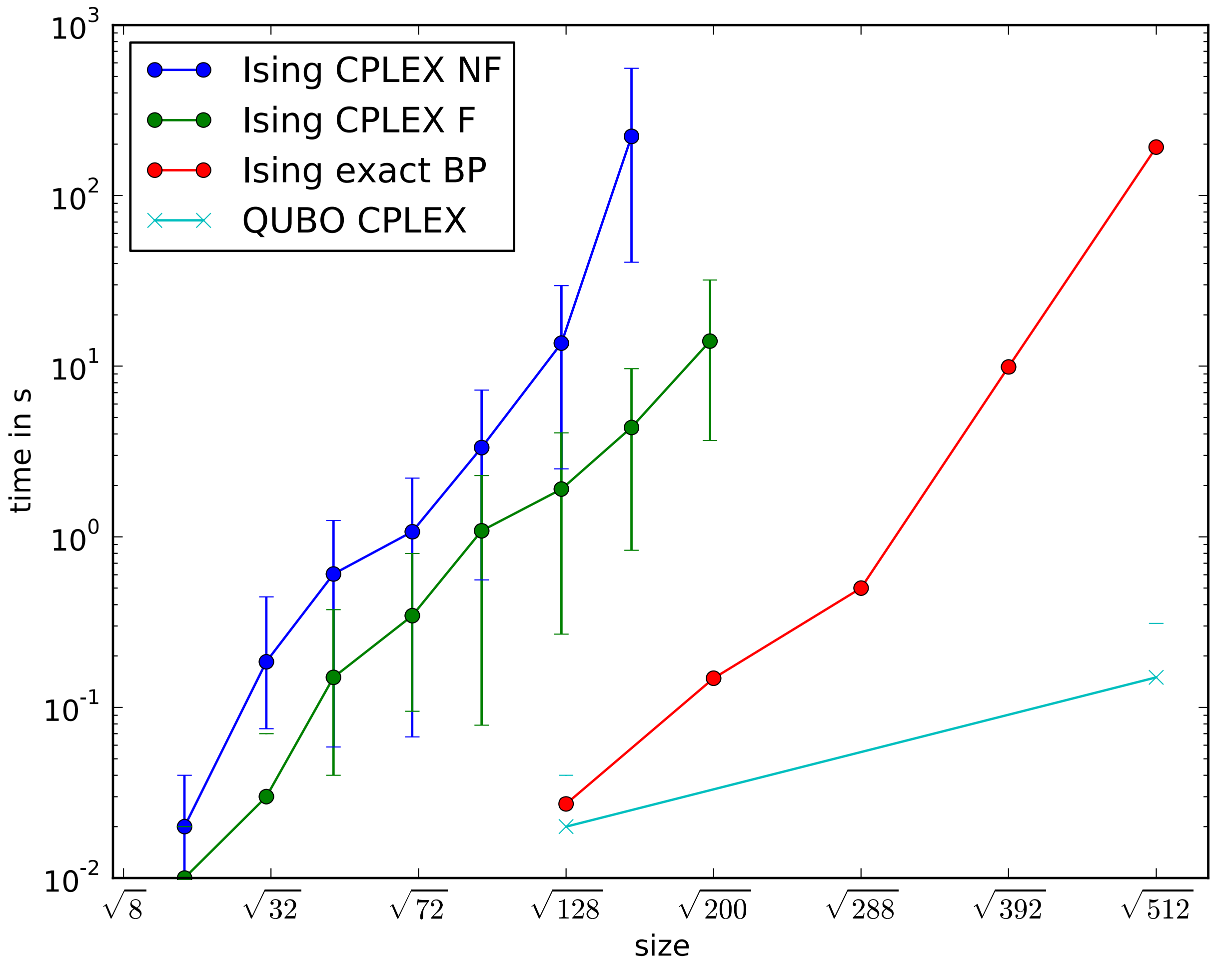

Quantum Annealing vs. Simulated Annealing

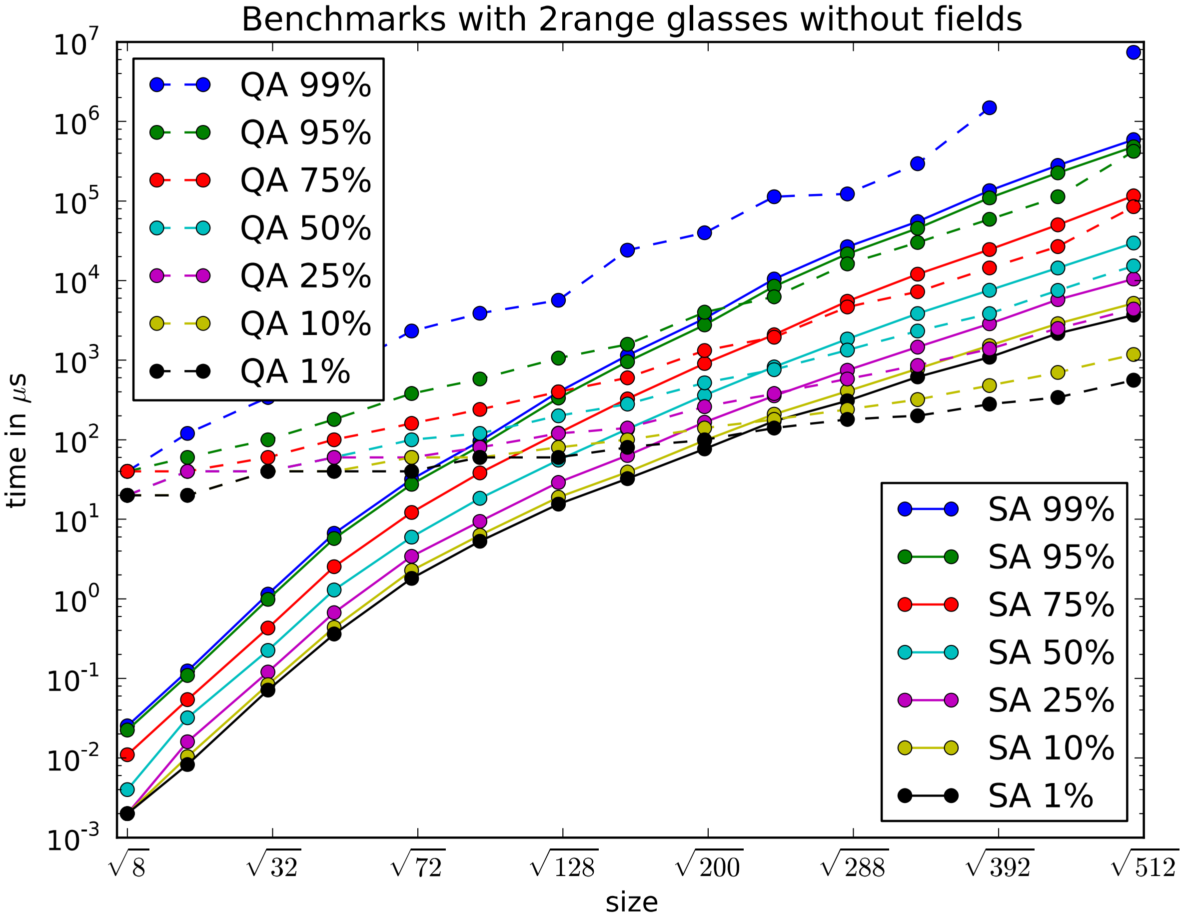

D-Wave 2 vs. Eight-core Pentium (USC & ETH)

Quantum Annealing vs. Simulated Annealing

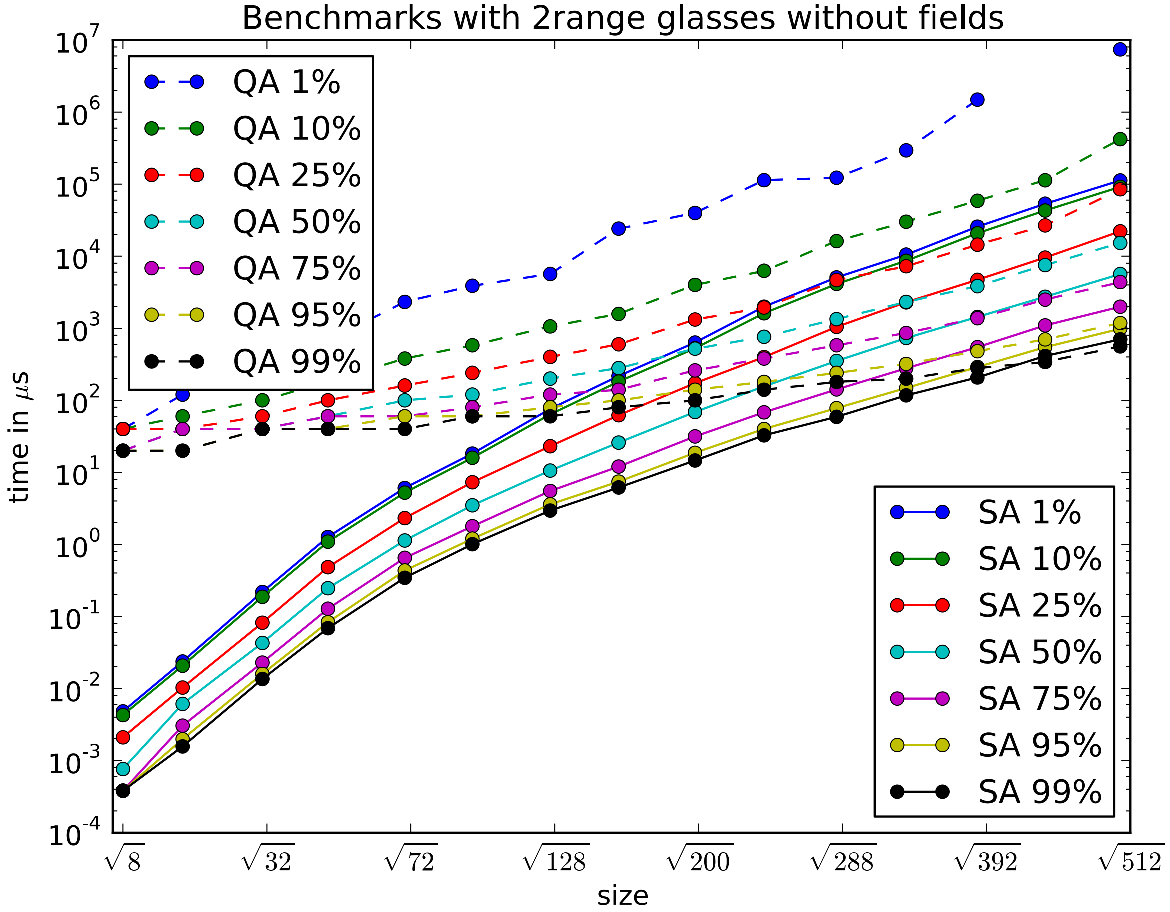

D-Wave 2 vs. Nvidia Kepler GPU (USC & ETH)

Hard problems motivated by satisfiability

Random walk to create frustrated loops that respect a planted solution

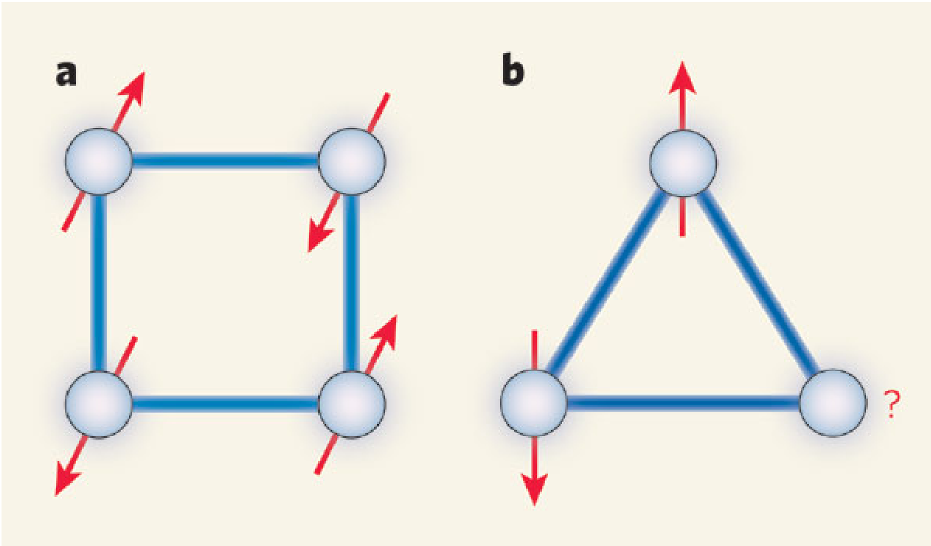

Frustration

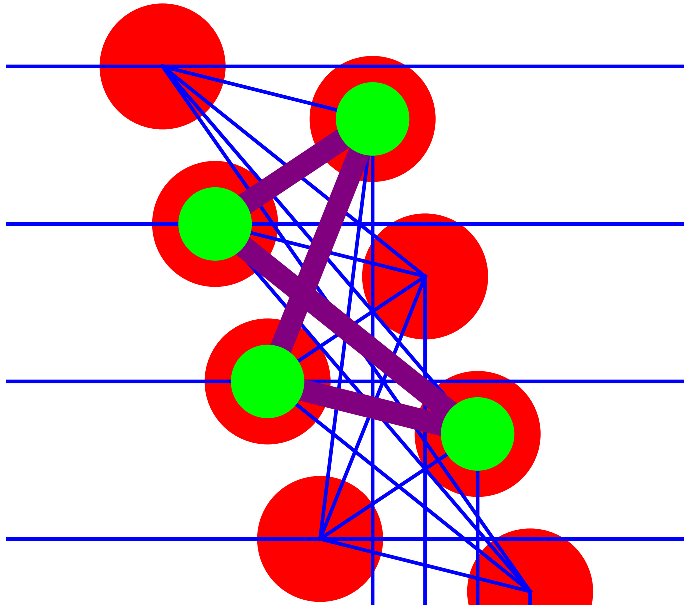

Itay Hen, Performance of D-Wave Two on Problems with Planted Solutions, AQC 2014

Hard problems for multiple heuristics

- "at least once with 99% chance"– a comparison.
- universal peak in hardness.

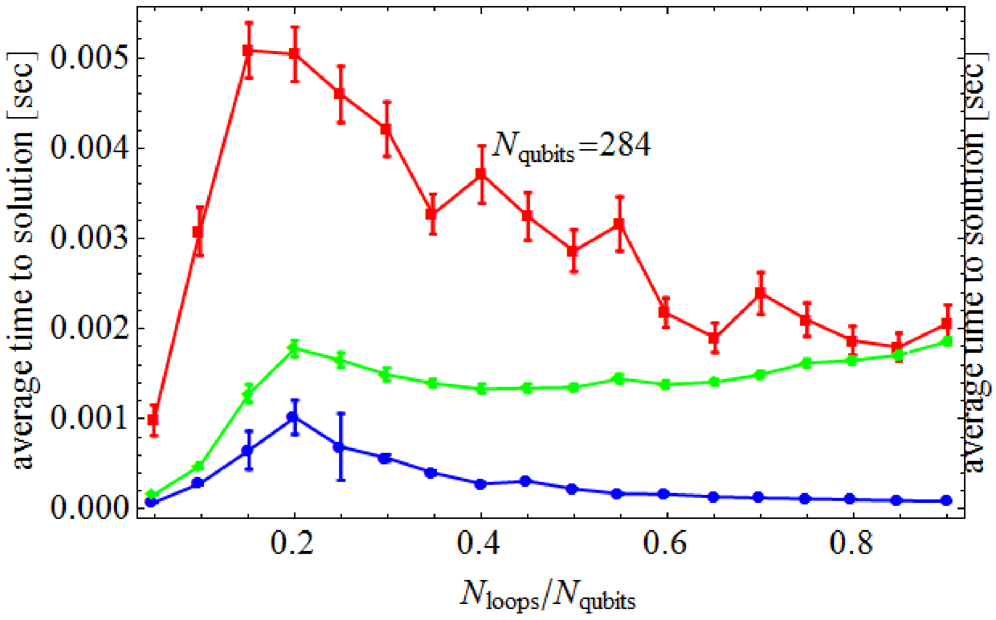

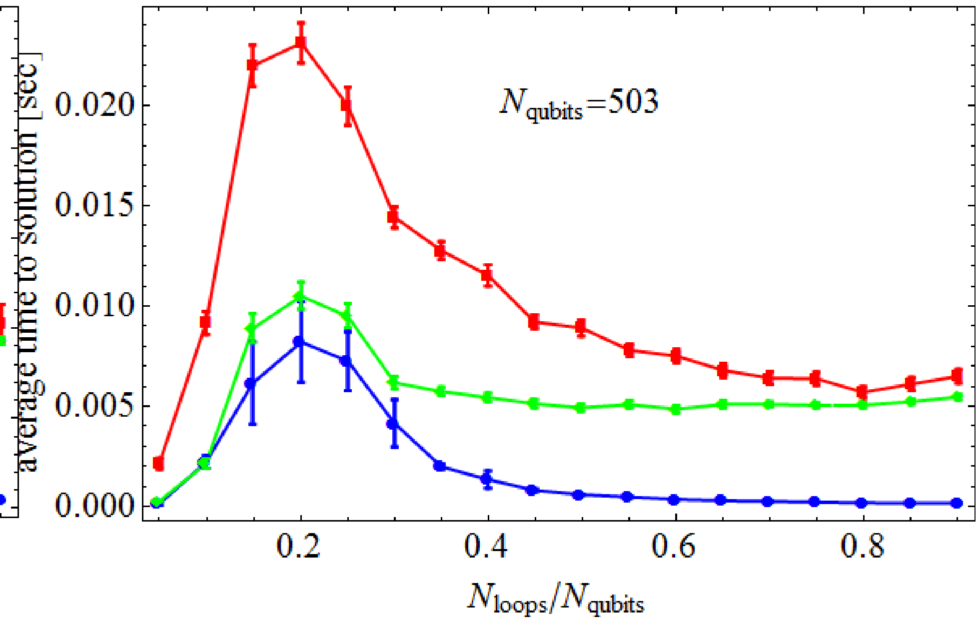

Performance on hard problems

- "average time to solution" scaling results.

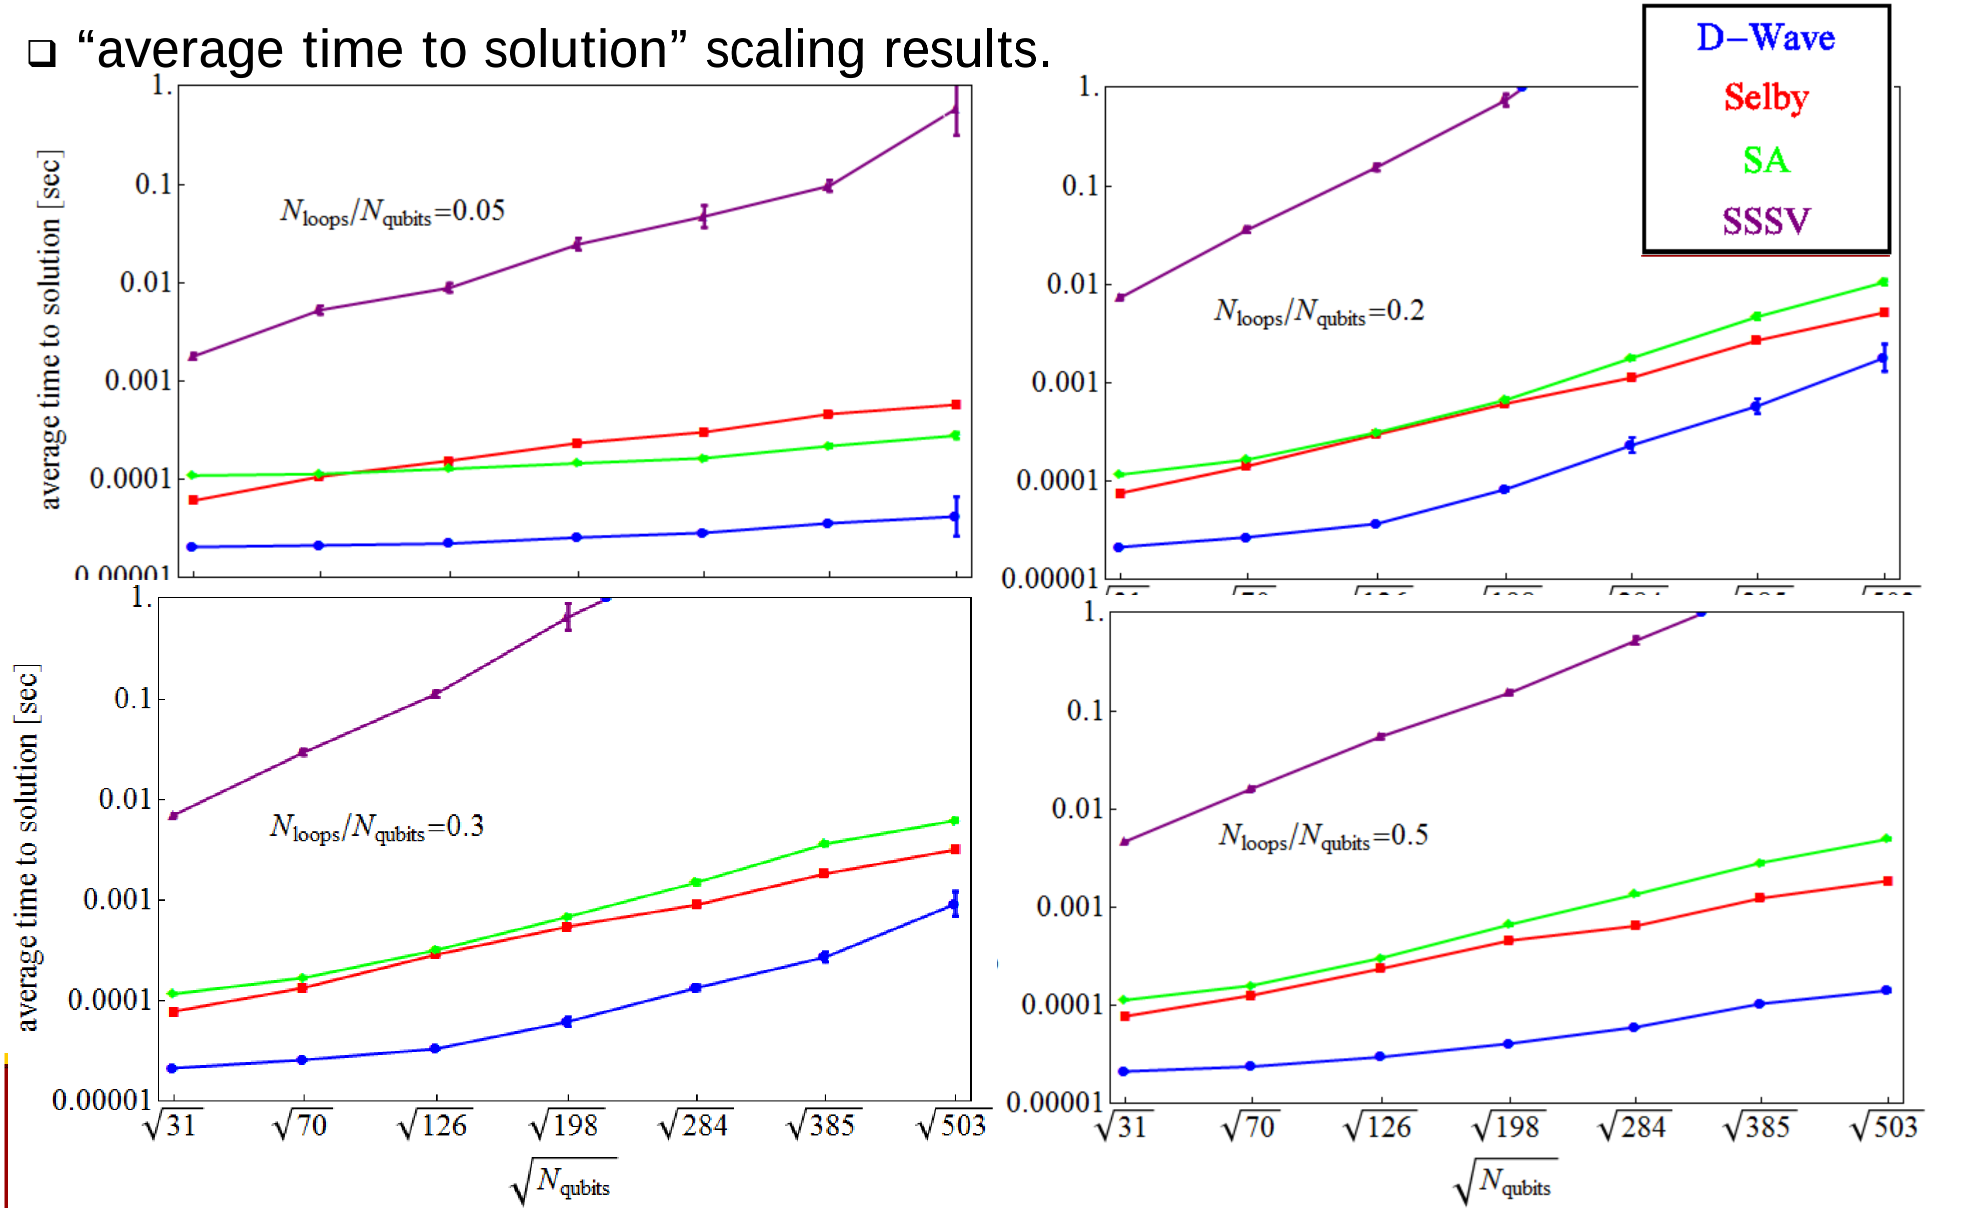

Speedup relative to what?

Better heuristics?

Selby motivated by earlier comparisons with simulated annealing

Multiple Xeon cores?

SAT solvers often don’t scale well

MPI overheads are O(us).

Not the only way to use CMOS

FPGA circuits (Victor Martin-Major’s earlier Janus talk)

ASICs ala Anton (Mark Moraes’s talk)

Full custom circuits

“D-Wave problem”

Best case for D-Wave

Not at all clear that this will extend to real applications

Outline

D-Wave Machine

Research Results to Date

Applications

Summary

.

What Problems Might be Amenable?

Problem	Application
Traveling salesman	Logistics, vehicle routing
Minimum Steiner tree	Circuit layout, network design
Graph coloring	Scheduling, register allocation
MAX-CLIQUE	Social networks, bioinformatics
QUBO	Machine learning
Integer Linear Programming	Natural language processing
Sub-graph isomorphism	Cheminformatics, drug discovery
Motion planning	Robotics
MAX-2SAT	Artificial intelligence

NP-complete problems from Wikipedia

On-Going Applications Research

Verification and Validation

Model checking for DARPA
Cyber physical system V&V at Lockheed Martin

Optimization of System Design

Collaboration with Lockheed Martin

Image Processing

Image registration at D-Wave
Image recognition at Google
Organic photovoltaic triage at Harvard

Missile Defense

Closely-spaced objects at Aerospace

Model Checking

- Given a:
  - Finite transition system *M*
  - A temporal property *p*
- The model checking problem:
  - Does *M* satisfy *p*?

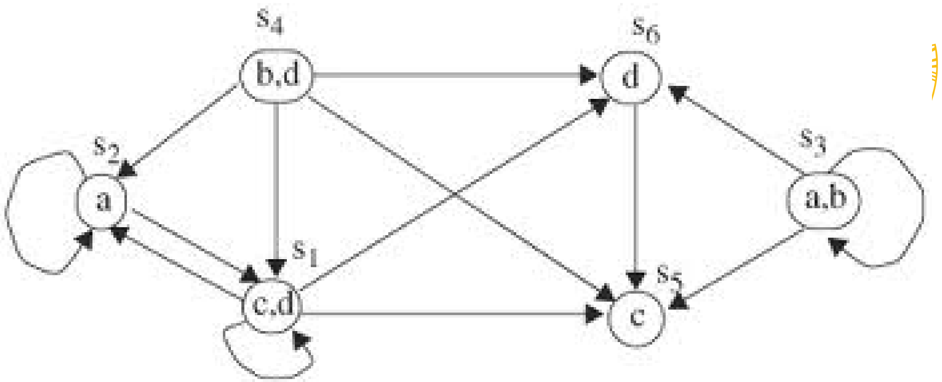

$$\Box((call \vee \Diamond open) \rightarrow ((\neg atfloor \vee \neg open)\ \mathcal{U}\ (open \vee ((atfloor \wedge \neg open)\ \mathcal{U}\ (open \vee ((\neg atfloor \wedge \neg open)\ \mathcal{U}\ (open \vee ((atfloor \wedge \neg open)\ \mathcal{U}\ (open \vee (\neg atfloor\ \mathcal{U}\ open))))))))))$$

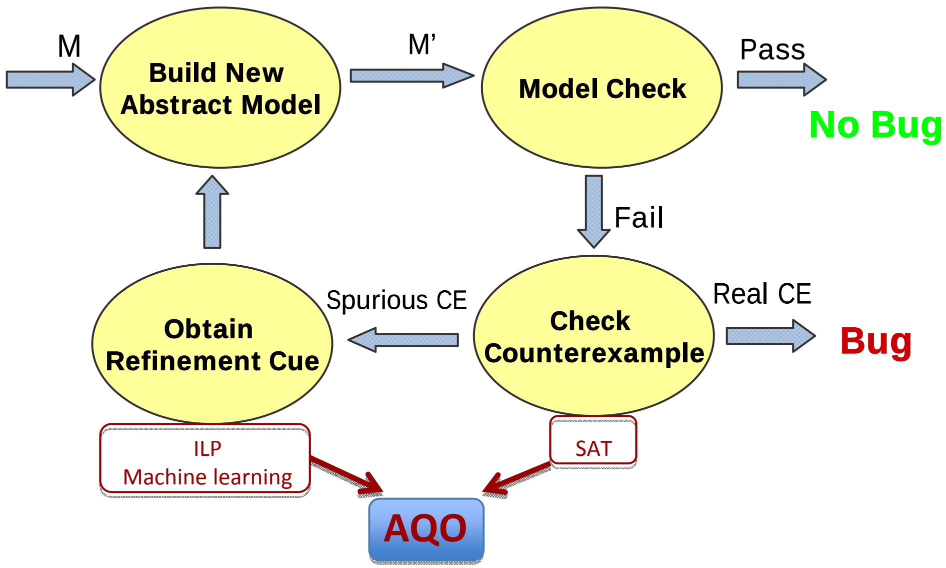

- Counter-example guided abstraction refinement
  - Identify combinatorial optimization components of classical algorithms and apply AQO

Nested Dissection Reordering

Best reordering algorithm for many linear systems

Widely used in practice today

Available in packages such as Metis and Scotch

Basic ND algorithm

Find a small separator to partition the matrix into two halves

- In general, this is an NP-complete problem

Refine the separator, "straightening" it

- minimize separator length
- balance the two subgraphs

Recursively partition the remaining subgraphs

Many Partitioning Heuristics Available

Ones I know used

Level Sets
George and Liu
Ashcraft and Grimes

Moment of Inertia

Fiduccia-Mattheyses

Spectral Bisection
Fiedler vectors

Ones I don't know of

Simulated Annealing

Quantum Annealing

Ising Models for NP-hard Problems

Graph bisection can be described as an Ising spin glass

J Phys **A19** 1605 (1986)

Since then, Ising models for many other problems have been found too

arXiv: 1302.5843

The general goal: find a cost function H (Hamiltonian) whose solutions correspond to the hard problem

Create energy penalties for suboptimal solutions

Ising model for Graph Partitioning

$$H = A \underbrace{\left(\sum_{i\in V} \sigma_i\right)^2}_{\text{penalize config. that doesn't bisect}} + B \sum_{ij\in E} \underbrace{\frac{1-\sigma_i\sigma_j}{2}}_{\text{penalty } B \text{ for each edge btwn two subsets}}$$

Partition the graph into two domains: spin up or down
Add energy penalty A to configurations that don't bisect
Add energy penalty B to edges in the separator
Finds an edge separator

Toy problem

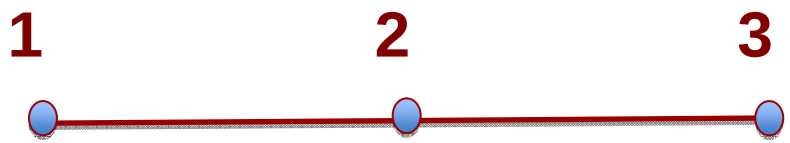

Solutions ^{T} * H * All possible solutions = Energy

Solutions^T

-	-	-
-	-	+
-	+	-
-	+	+
+	-	-
+	-	+
+	+	-
+	+	+

*

H

1	0	1
0	1	0
1	0	1

*

All possible solutions

-	-	-	-	+	+	+	+
-	-	+	+	-	-	+	+
-	+	-	+	-	+	-	+

=

Energy

5
1
5
1
1
5
1
5

Direct embedding of Ising Model to Chimeral Graph

Embedded 10-Node Line

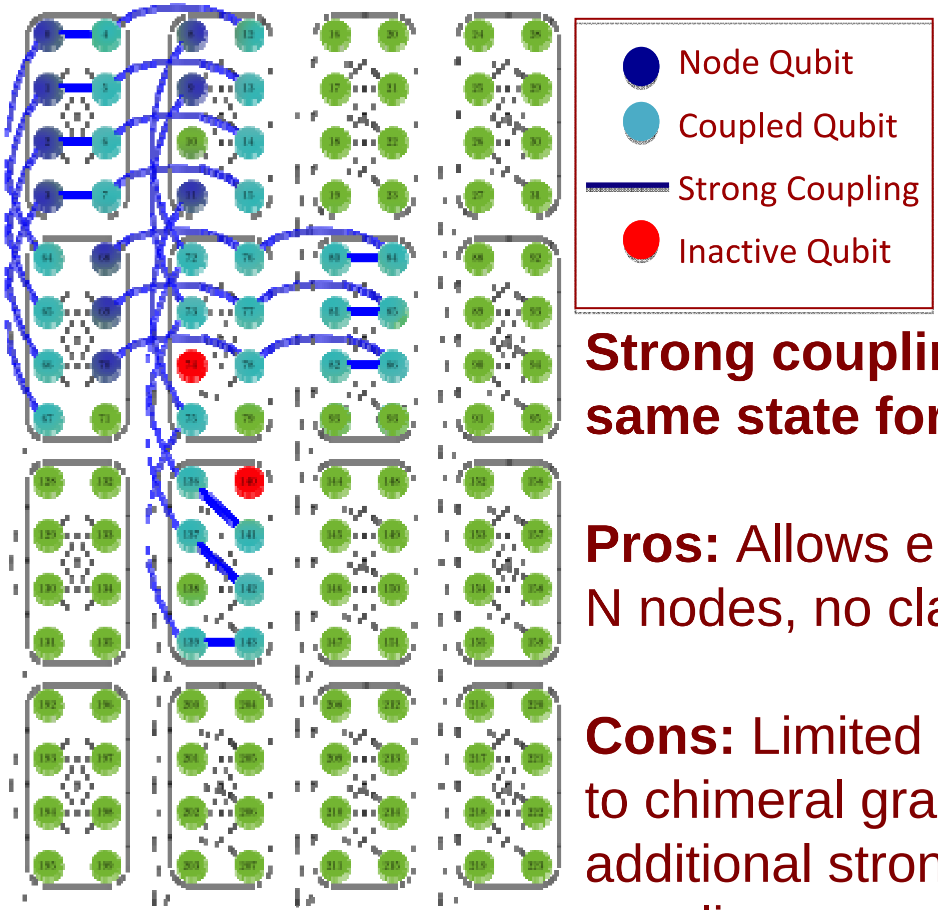

Use strong couplings between qubits to create virtual nodes with full connectivity

Strong coupling forces coupled qubits into the same state for low energy solutions

Pros: Allows embedding of any graph topology with N nodes, no classical post-processing necessary

Cons: Limited by number of qubits, Mapping problem to chimeral graph for large N is NP-Hard, Each additional strong coupling reduces effect of weaker couplings

Test Results: Fully Embedded 10 Node Line Graph

Ideal Partition appears as the lowest energy solution state

E (-30)	N	State
-7.40E-03	**11**	**[1 1 1 1 1 -1 -1 -1 -1 -1]**
-7.40E-03	**1**	**[-1 -1 -1 -1 -1 1 1 1 1 1]**
-6.40E-03	29	[1 1 1 1 -1 -1 -1 -1 -1 -1]
-6.40E-03	3	[1 1 1 1 1 1 -1 -1 -1 -1]
-6.40E-03	1	[-1 -1 -1 -1 1 1 1 1 1 1]
-6.40E-03	3	[-1 -1 -1 -1 -1 -1 1 1 1 1]
-6.00E-03	35	[-1 -1 1 1 1 1 1 -1 -1 -1]
-6.00E-03	32	[-1 1 1 1 1 1 -1 -1 -1 -1]

Dense Hamiltonian Coupling Factors:

Strong: **-1**
Adjacent: **.0002**
Weak: **.0005**

Small differences in solution state energies due to large relative weight of strong couplings

Future Quantum Annealers with higher connectivity and more qubits may make Direct Embedding more realistic

Outline

D-Wave Machine

Research Results to Date

Applications

Summary

.

Summary

After little over a decade, adiabatic quantum annealing is moving from theory to practice.

Today's D-Wave system raises a variety of research questions that USC, Lockheed Martin, and our colleagues are jointly investigating

Future D-Wave systems might soon be the most powerful on Earth for a range of problems

What those might be, and how to program them, are still open research problems.

Open Research Problems

Why does the D-Wave even work?

Its an open system

How much quantum speedup will there be?

Any? If so, on what problems?

What applications will it ultimately solve?

We’ve had half a century to find competing heuristics

How should you program it?

Specifically excluded from recent research programs

What should the topology be?

Reduce critical scaling limitation

Other adiabatic quantum systems will face these.

These questions are all bigger than just D-Wave

Back Up

Case Study: 4-city Traveling Salesman (LM 2011)

- 4 city directed Traveling Salesman Problem
- 12 logical qubits (1 per directed link)
  - Embeddable on 92-qubit processor
- Found optimal solution in 10% of iterations

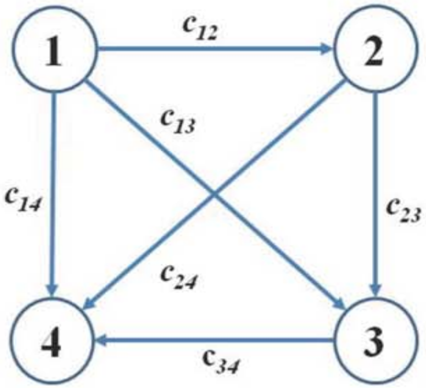

Link Costs

c12	5.86
c13	9.70
c14	5.51
c21	6.75
c23	3.91
c24	8.18
c31	1.31
c32	1.31
c34	4.39
c41	5.86
c42	3.30
c43	5.51

Constraints

	x12	x13	x14	x21	x23	x24	x31	x32	x34	x41	x42	x43
x12	-2	2	2	1	0	0	0	2	0	0	2	0
x13	0	-2	2	0	2	0	1	0	0	0	0	2
x14	0	0	-2	0	0	2	0	0	2	1	0	0
x21	0	0	0	-2	2	2	2	0	0	2	0	0
x23	0	0	0	0	-2	2	0	1	0	0	0	2
x24	0	0	0	0	0	-2	0	0	2	0	1	0
x31	0	0	0	0	0	0	-2	2	2	2	0	0
x32	0	0	0	0	0	0	0	-2	2	0	2	0
x34	0	0	0	0	0	0	0	0	-2	0	0	1
x41	0	0	0	0	0	0	0	0	0	-2	2	2
x42	0	0	0	0	0	0	0	0	0	0	-2	2
x43	0	0	0	0	0	0	0	0	0	0	0	-2

Energy distribution of annealing iterations

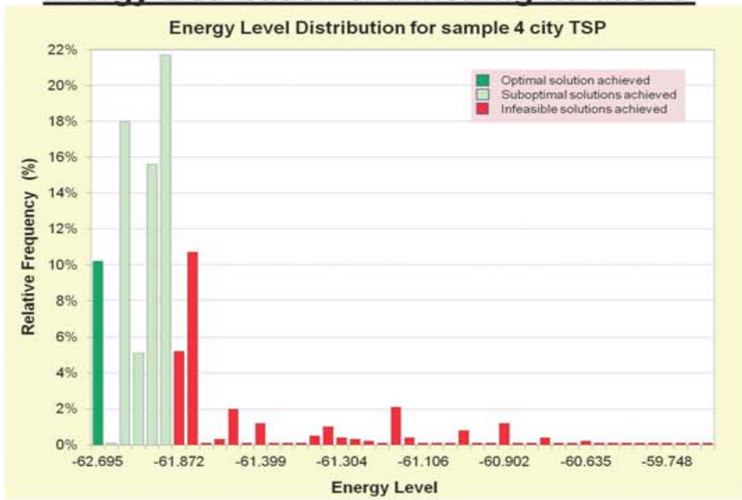

Slide Courtesy of Steve Adachi, Lockheed Martin

Case Study: Protein Folding (Harvard/D-Wave)

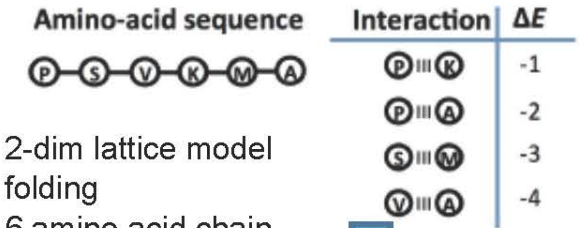

- Simplified 2-dim lattice model of protein folding
- Modeled a 6 amino acid chain

- Model has 18 logical qubits
- Embedding took 81 physical qubits

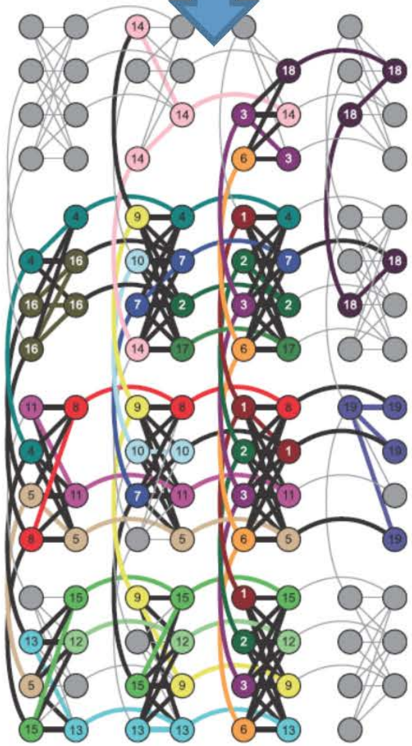

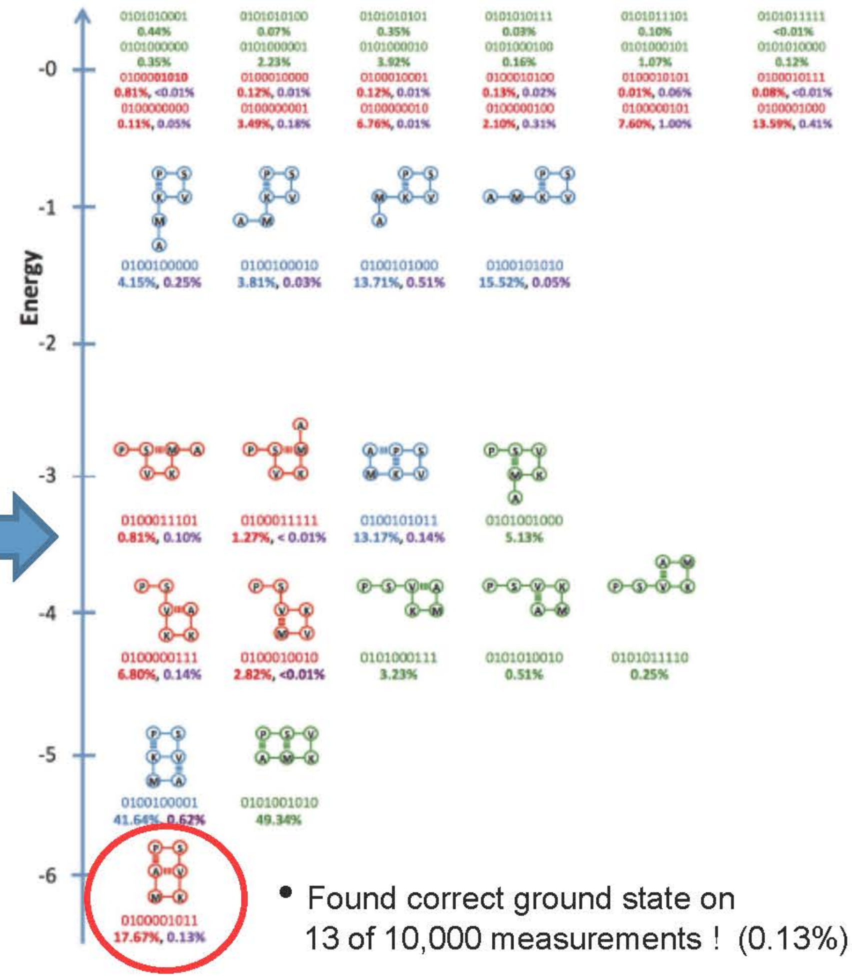

- Found correct ground state on 13 of 10,000 measurements ! (0.13%)

Ref: A. Perdomo-Ortiz, N. Dickson, M. Drew-Brook, G. Rose & A. Aspuru-Guzik, "Finding low-energy conformations of lattice protein models by quantum annealing," doi:10.1038/srep00571.

Slide Courtesy of Steve Adachi, Lockheed Martin