Dawn of the Era of Quantum Computation

Bob Lucas
USC – Lockheed Martin Quantum Computing Center
August 5, 2013

Outline

Historical Perspective

Tomorrow's Uncertainty

Adiabatic Quantum Annealing

Summary

Long History of DOE Leadership in HPC

John von Neumann
Head of AEC

Los Alamos IBM Stretch
Introduced pipelines

Los Alamos Cray 1
Vector mainframe

Oak Ridge Intel iPSC 1
Commercial parallel computer

Sandia ASCI Red
Aggressive COTS-based scaling

Today's Leadership Computing Facilities
Continued tradition of National leadership

Long History of DOE Leadership in HPC

Scientific Disciplines	Math and CS
Physics	Differential Equations
Chemistry	Sparse Solvers
Biology	System Software
Material Science	Programming Environments
Earth Sciences	Visualization

Looming Challenge Thirty Years Ago

Technology driver transition

LSI to VLSI
Emergence of MOS, RISC microprocessors
Availability of dense semiconductor DRAM
Local area and system area networks

Paradigm Shift required

Exploit "killer micro"
Increase performance, reduce cost
Transition away from vector, PVP, and SIMD-array
How to structure and program new class of systems

Everything was about to change

How do we program future class of systems
Do we keep a shared memory model
New programming methods; what about legacy codes

The Future was Uncertain

BBN TC-2000 NUMA

Denelcor HEP multithreaded

Intel Touchstone Delta

TMC CM-2

Alliant FX/8

KSR-1

Cray T3D/E

J-Machine

Ncube

Cray C90

UIUC/CSRD Cedar

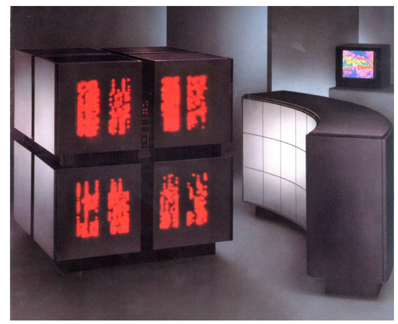

HPC Finally Reached a Stable Point

COTS Components

Leverage mainstream commercial processors

Cache-based hierarchical memories

Distributed Memory Architecture

Familiar software on processors (Fortran, C, and C++)

Message Passing Interface (MPI)

Evolving for two decades

Beowulf is the epitome

All off-the-self

PC components

Ethernet networks

Linux

gcc

DOE Made Major Contributions

Programming model

Message Passing Interface

ANL's MPICH was the reference implementation

OpenMPI & OpenMP today

Performance analysis tools: Jumpshot, TAU, HPCToolkit, etc.

Mathematics

New scalable algorithms (e.g., AMG)

Reference implementations: PETSc, HYPRE, BOXLIB, etc.

System Software

Micro-kernel O/S

Parallel file systems

Performance modeling

Outline

Historical Perspective

Tomorrow's Uncertainty

Adiabatic Quantum Annealing

Summary

The End of Dennard Scaling

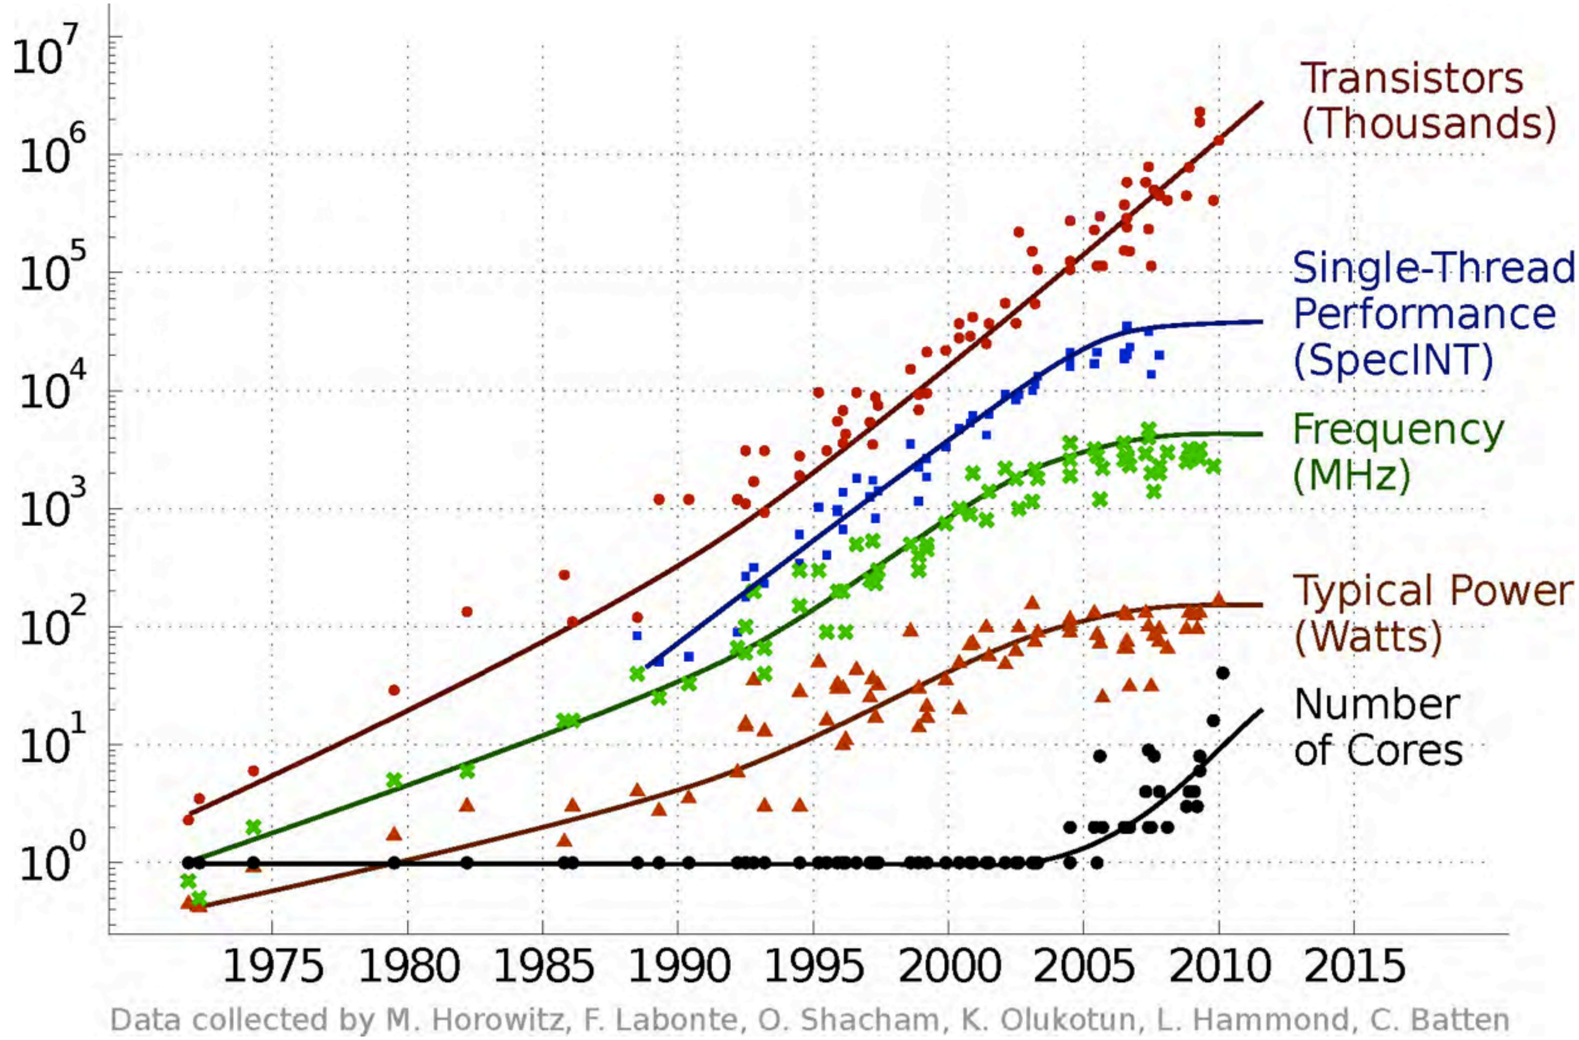

Uncertain Future has Returned

End of Moore's Law looming

Already lost relevance to strong-scaled problems

New Technologies on the near horizon

Phase change memories

Stacked dies

Exponential growth in parallelism

Tianhe-2 has 3M cores

Heterogeneous systems increase Flops/watt

Multicore chips have partitioned shared memory

Power is a constraint

Resilience a very real concern

Need More Capability?

Exploit a New Phenomenon
D-Wave Quantum Annealer

Massive Scaling
Tianhe-2 (3M cores)

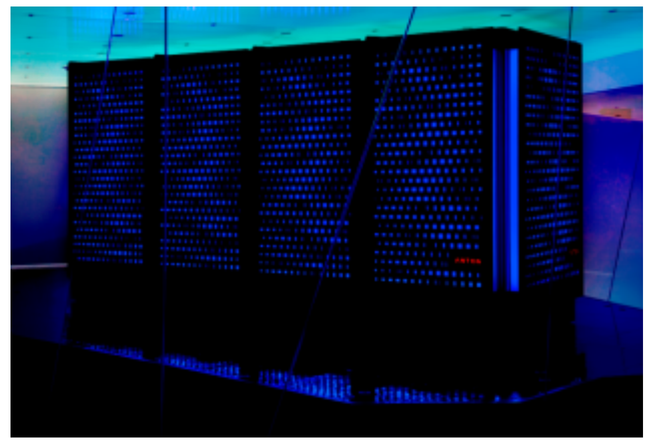

Application Specific Systems
D.E. Shaw Research Anton

Exascale Challenges

Energy
Hundreds of MWatts

Concurrency
Billions of ALUs

Memory
DRAM falling off Moore's Law

Resilience
Soft error rates expected to climb

Algorithms
Dot products and Courant numbers will be huge Amdahl fractions

Specialized Systems

Nothing new

Enigma

Germans invent machine cryptography

Bombe

Poles invent machine cryptanalysis

Specialized Systems

Not new to science either

Many QCD devices

Commercial protein folding too

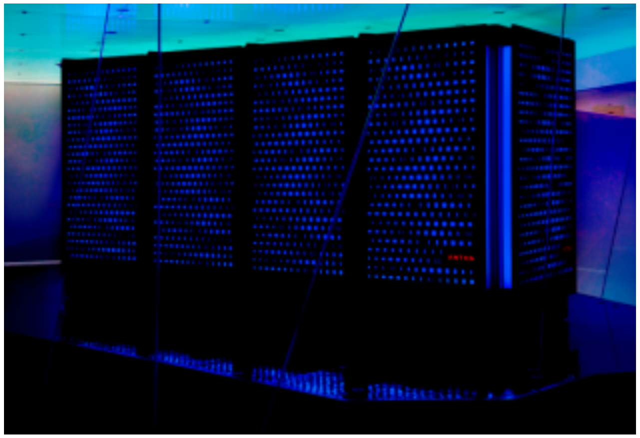

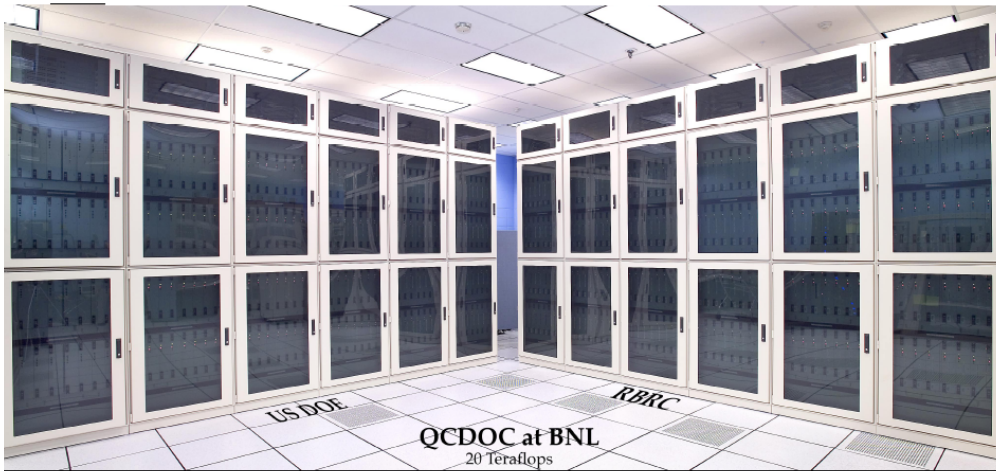

Outline

Historical Perspective

Tomorrow's Uncertainty

Adiabatic Quantum Annealing

Summary

Predicting the Future

Rolf Landauer
"Information is physics"

Memo to IBM

**The transistor:
Nothing to worry about …**

Predicting the Future is Fraught with Uncertainty

Rolf Landauer

**The memo was precisely right about the first transistor…
but not the second.**

Look Beyond CMOS

D-Wave System Collage

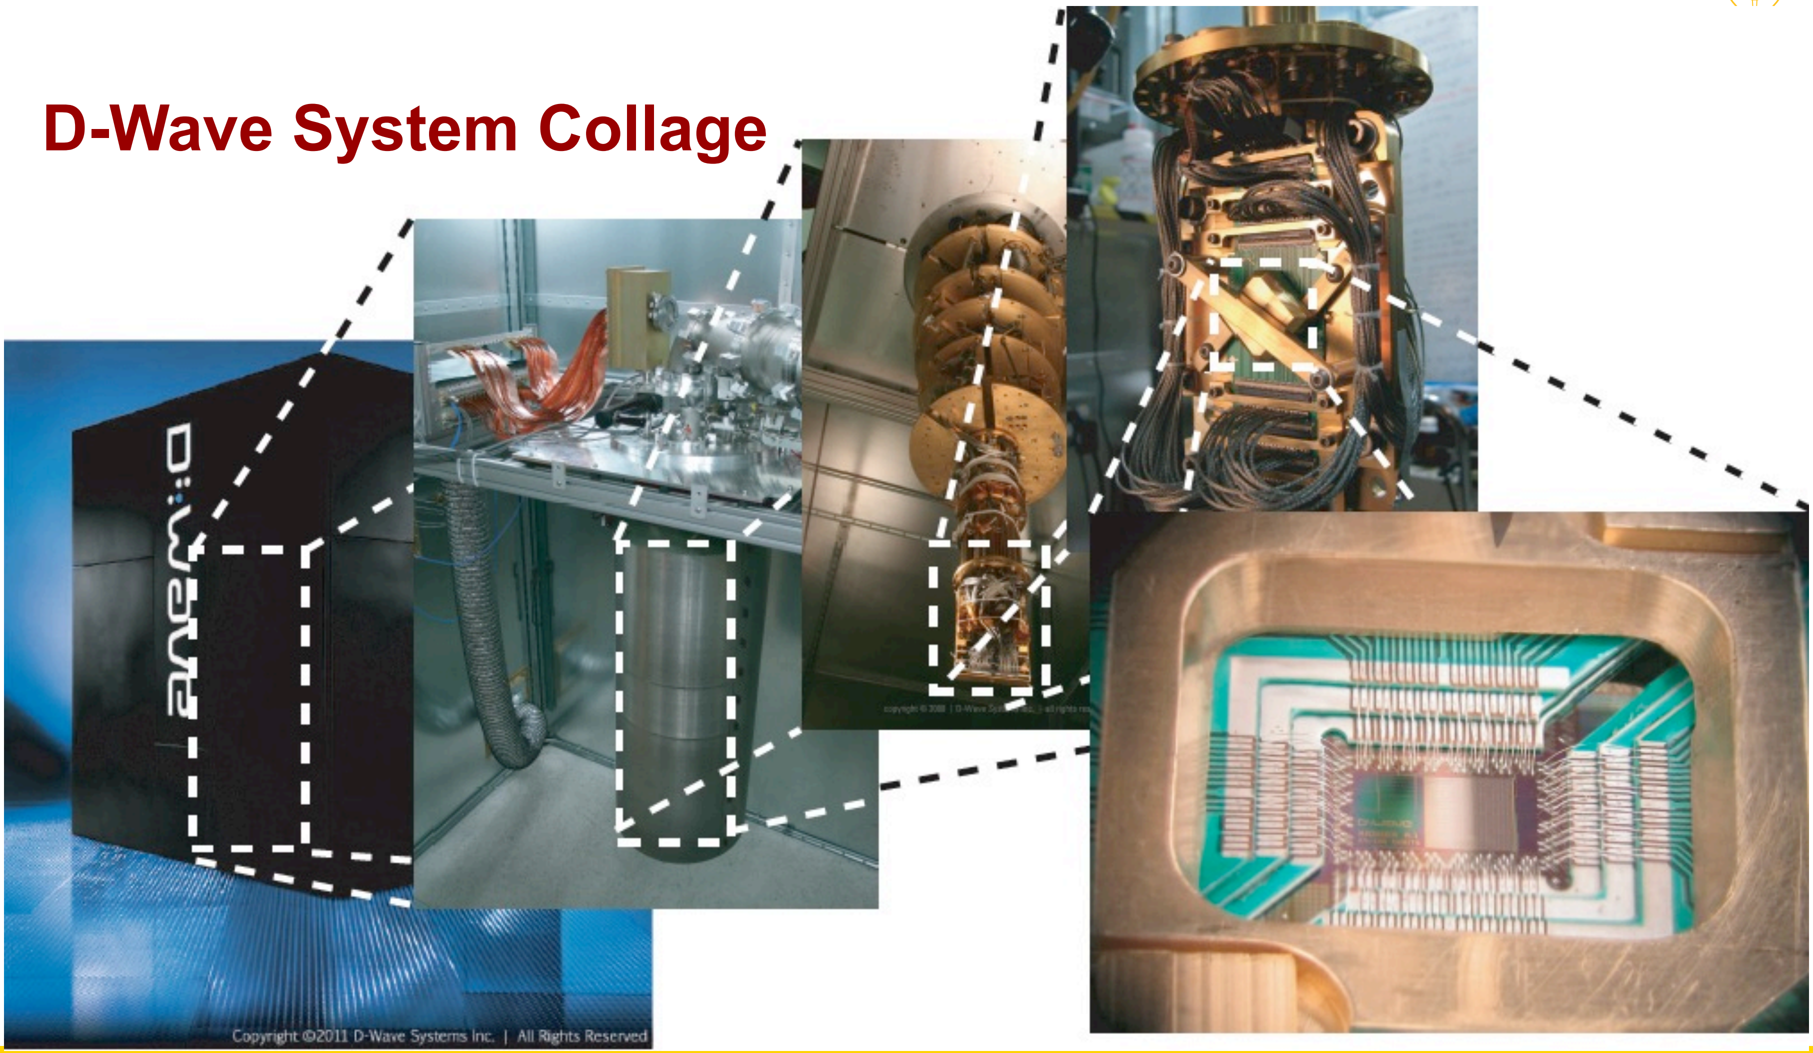

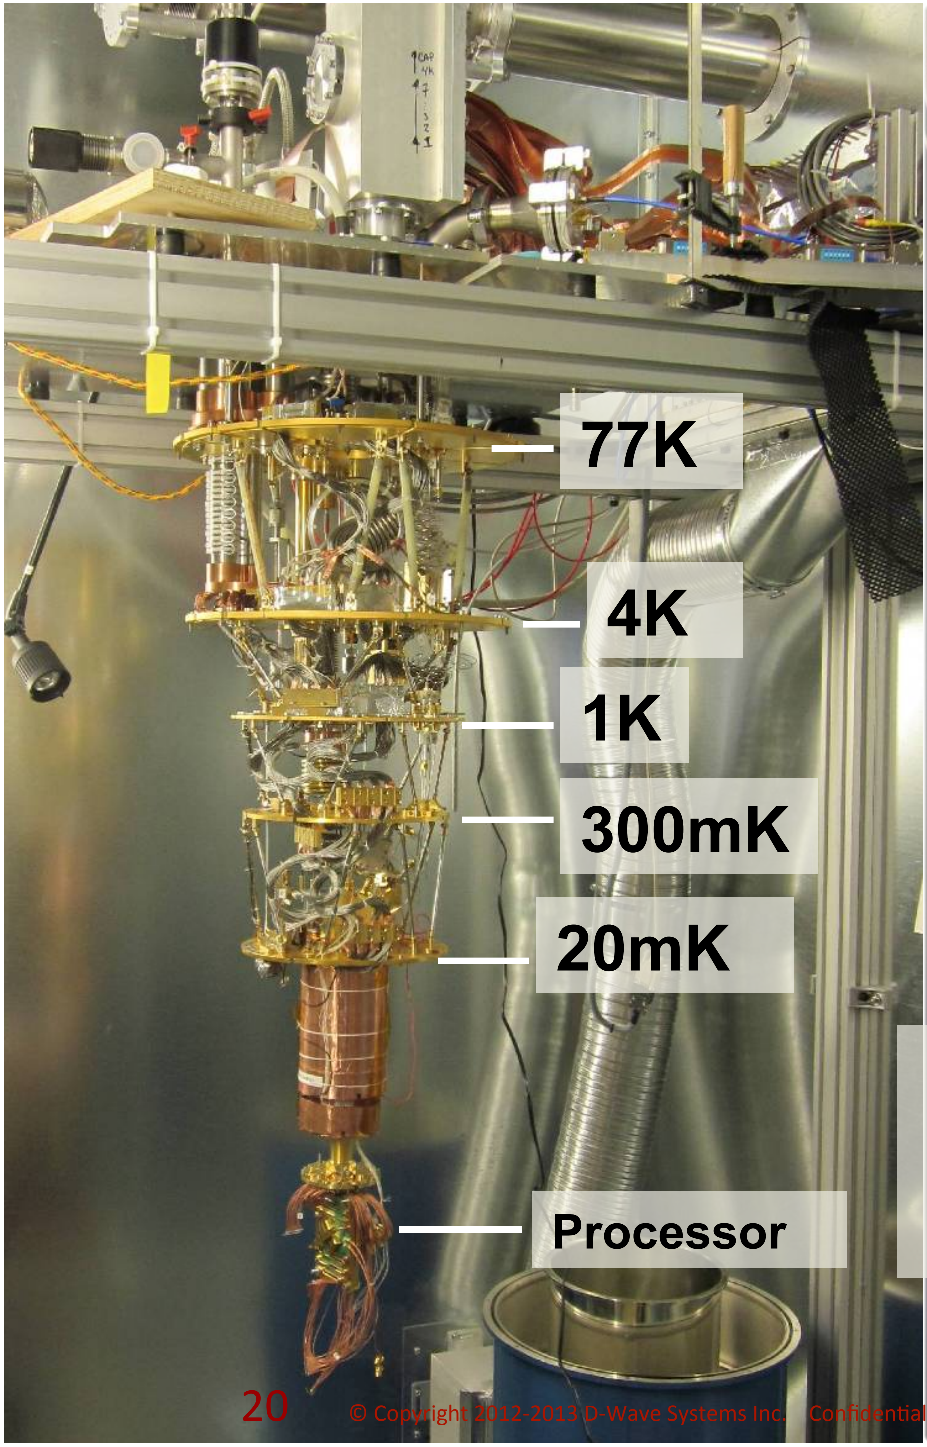
77K
4K
1K
300mK
20mK
Processor

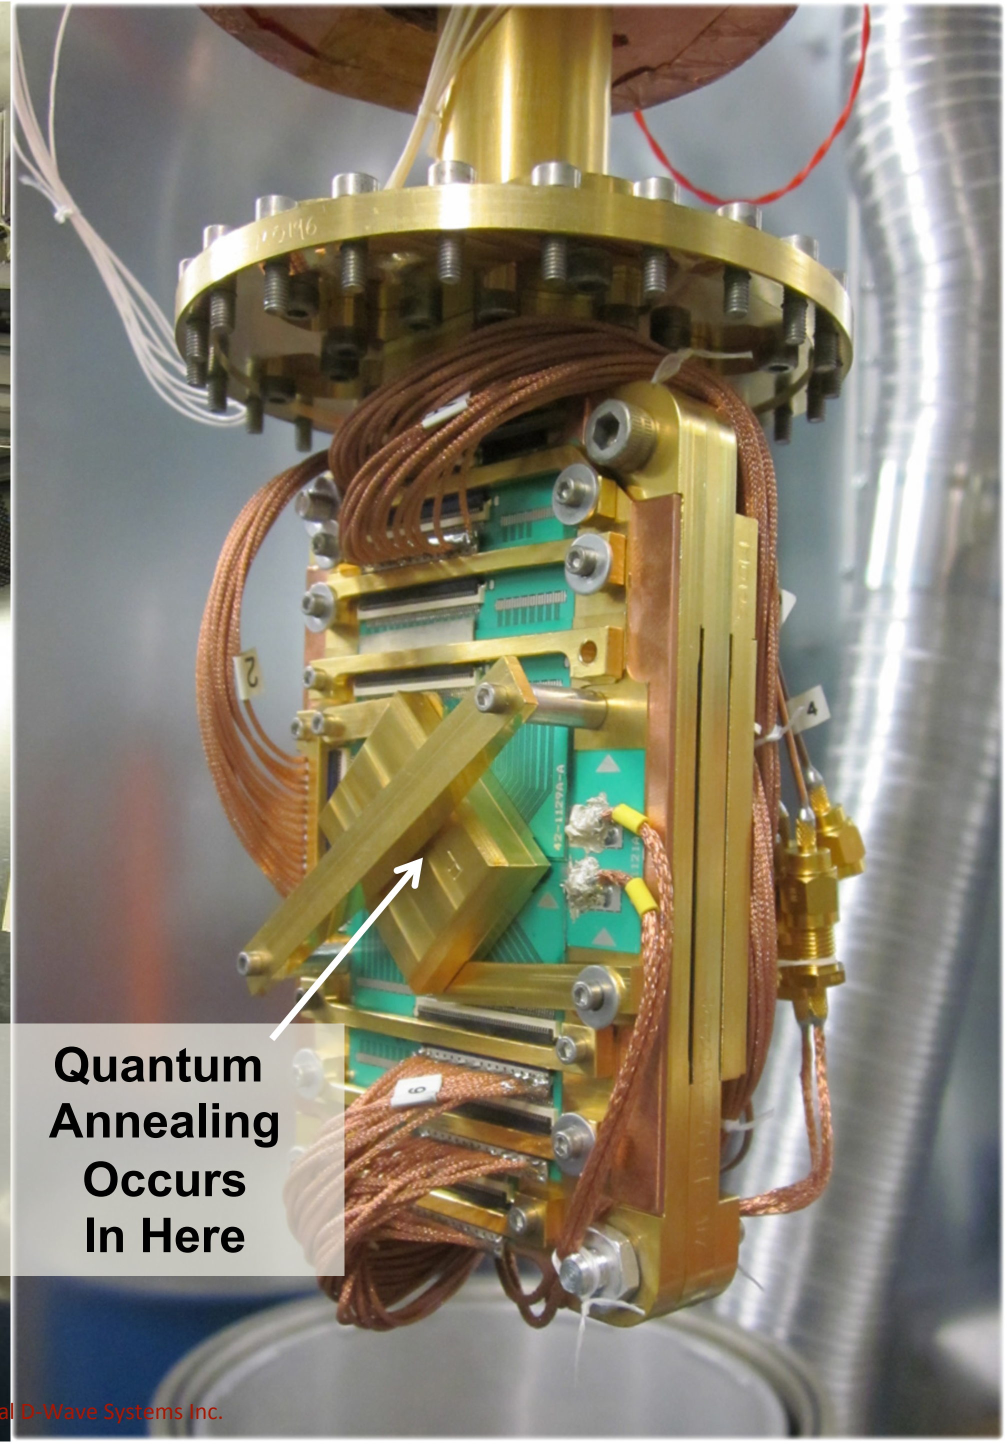
Quantum
Annealing
Occurs
In Here

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)

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

$t \in [0, t_f]$

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

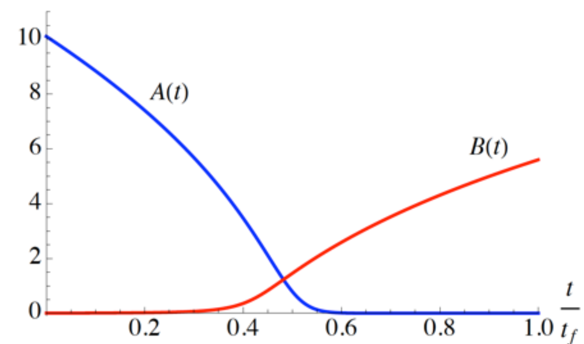

Graph Embedding implemented on DW-1 via Chimera graph retains NP-hardness (V. Choi, 2010)

D-Wave Eight Qubit Unit Cell

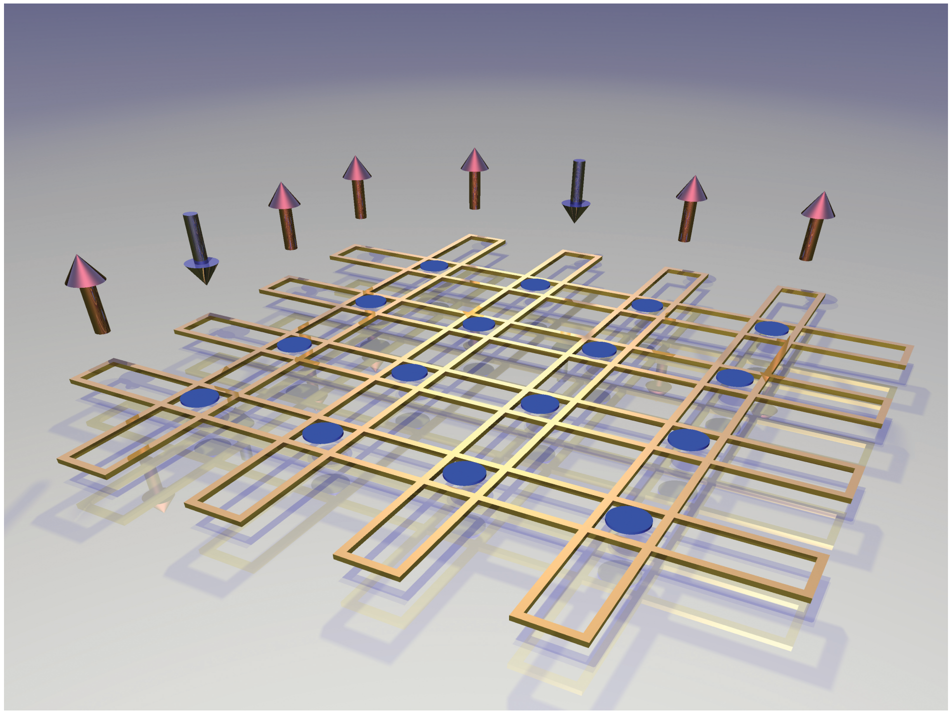

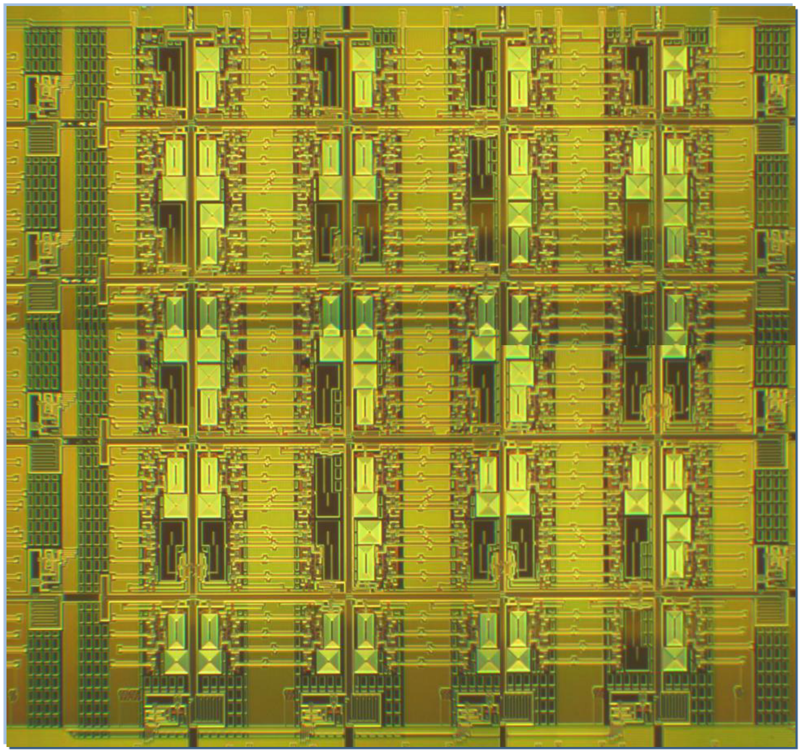

Images courtesy D-Wave

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.

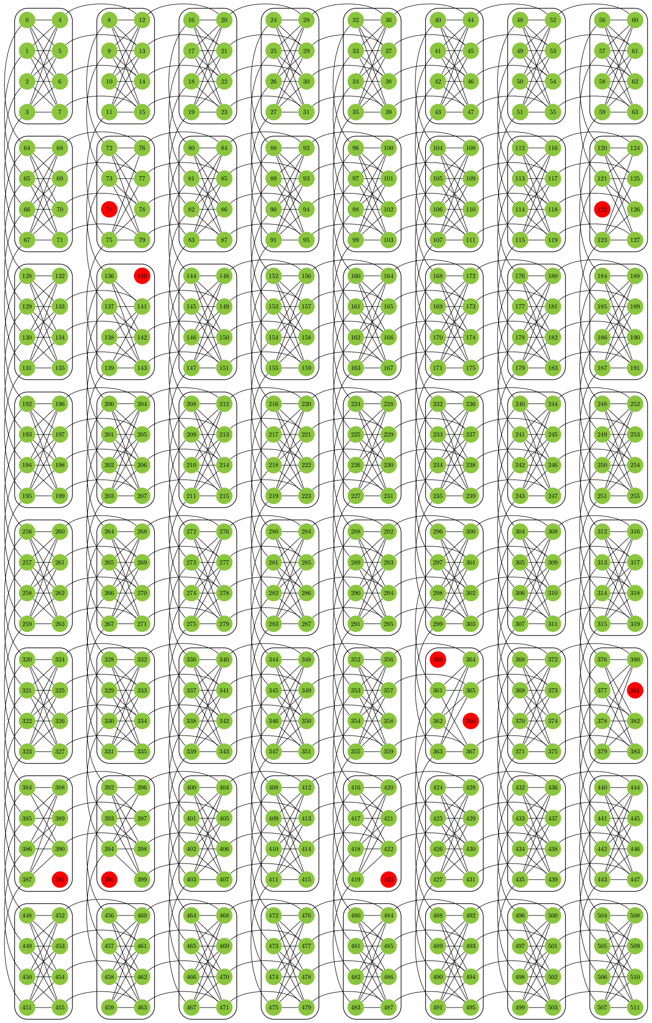

D-Wave's Version of "Moore's Law"

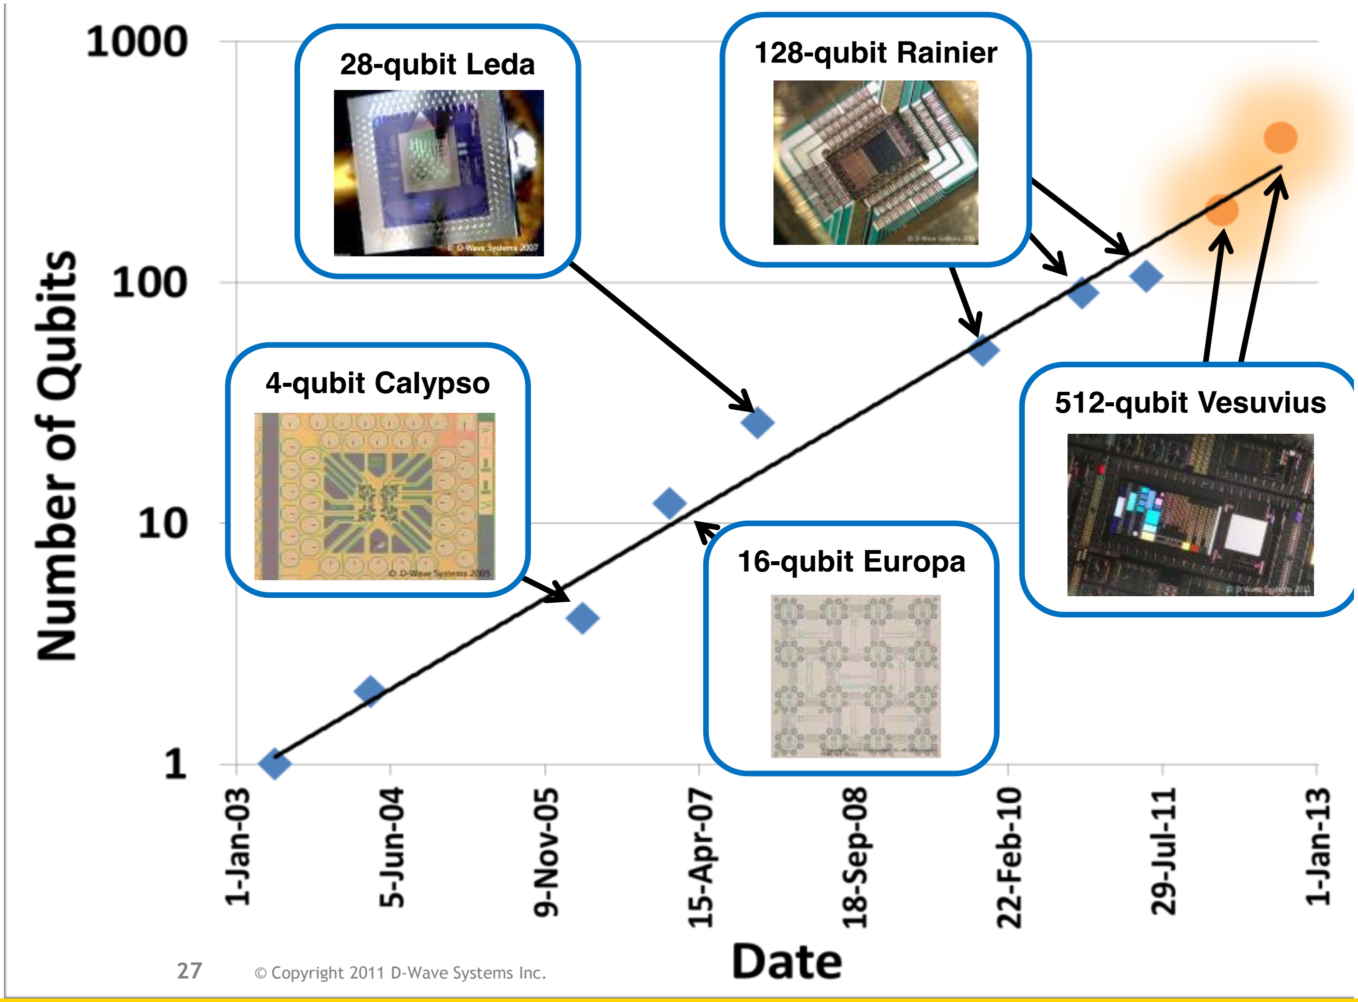

Quantum Signature: Degenerate Ising Hamiltonian

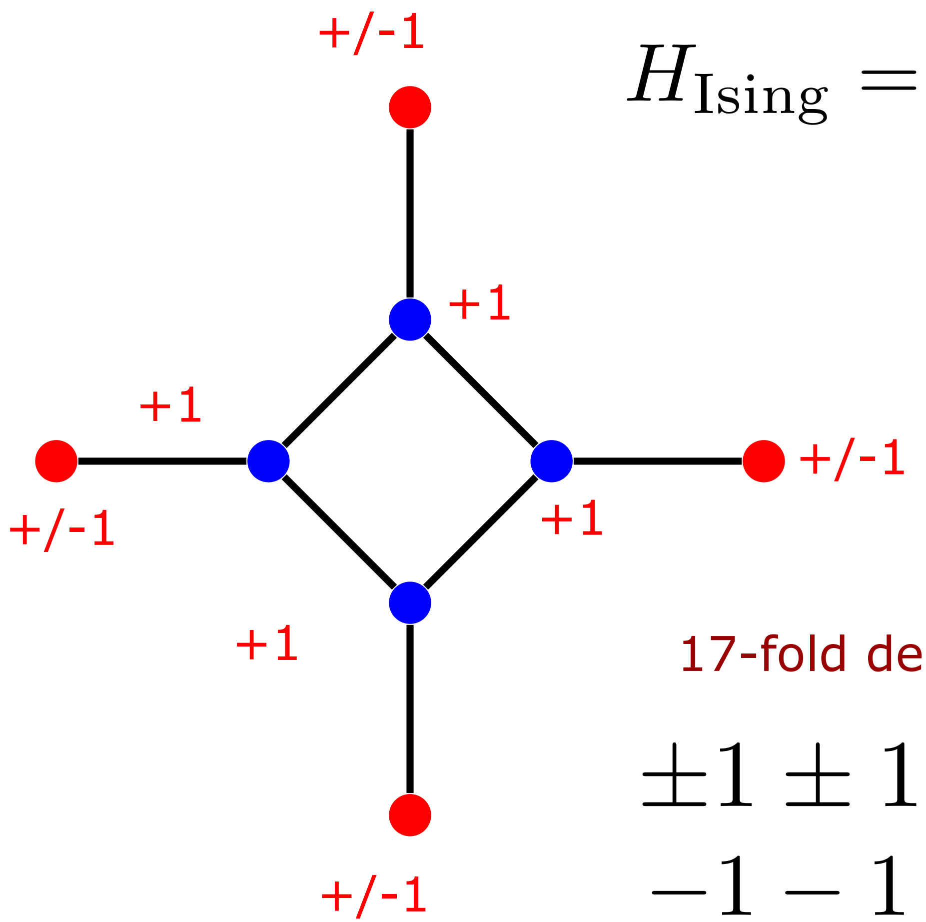

$$H_{\text{Ising}} = \sum_j h_j \sigma_j^z + \sum_{(j,k)\in E} J_{jk}\sigma_j^z \sigma_k^z$$

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

17-fold degenerate ground space:

$$\pm1 \pm1 \pm1 \pm1 \quad 1 \quad 1 \quad 1 \quad 1$$
$$-1 -1 -1 -1 -1 -1 -1 -1$$

Simulated Annealing At Several Speeds

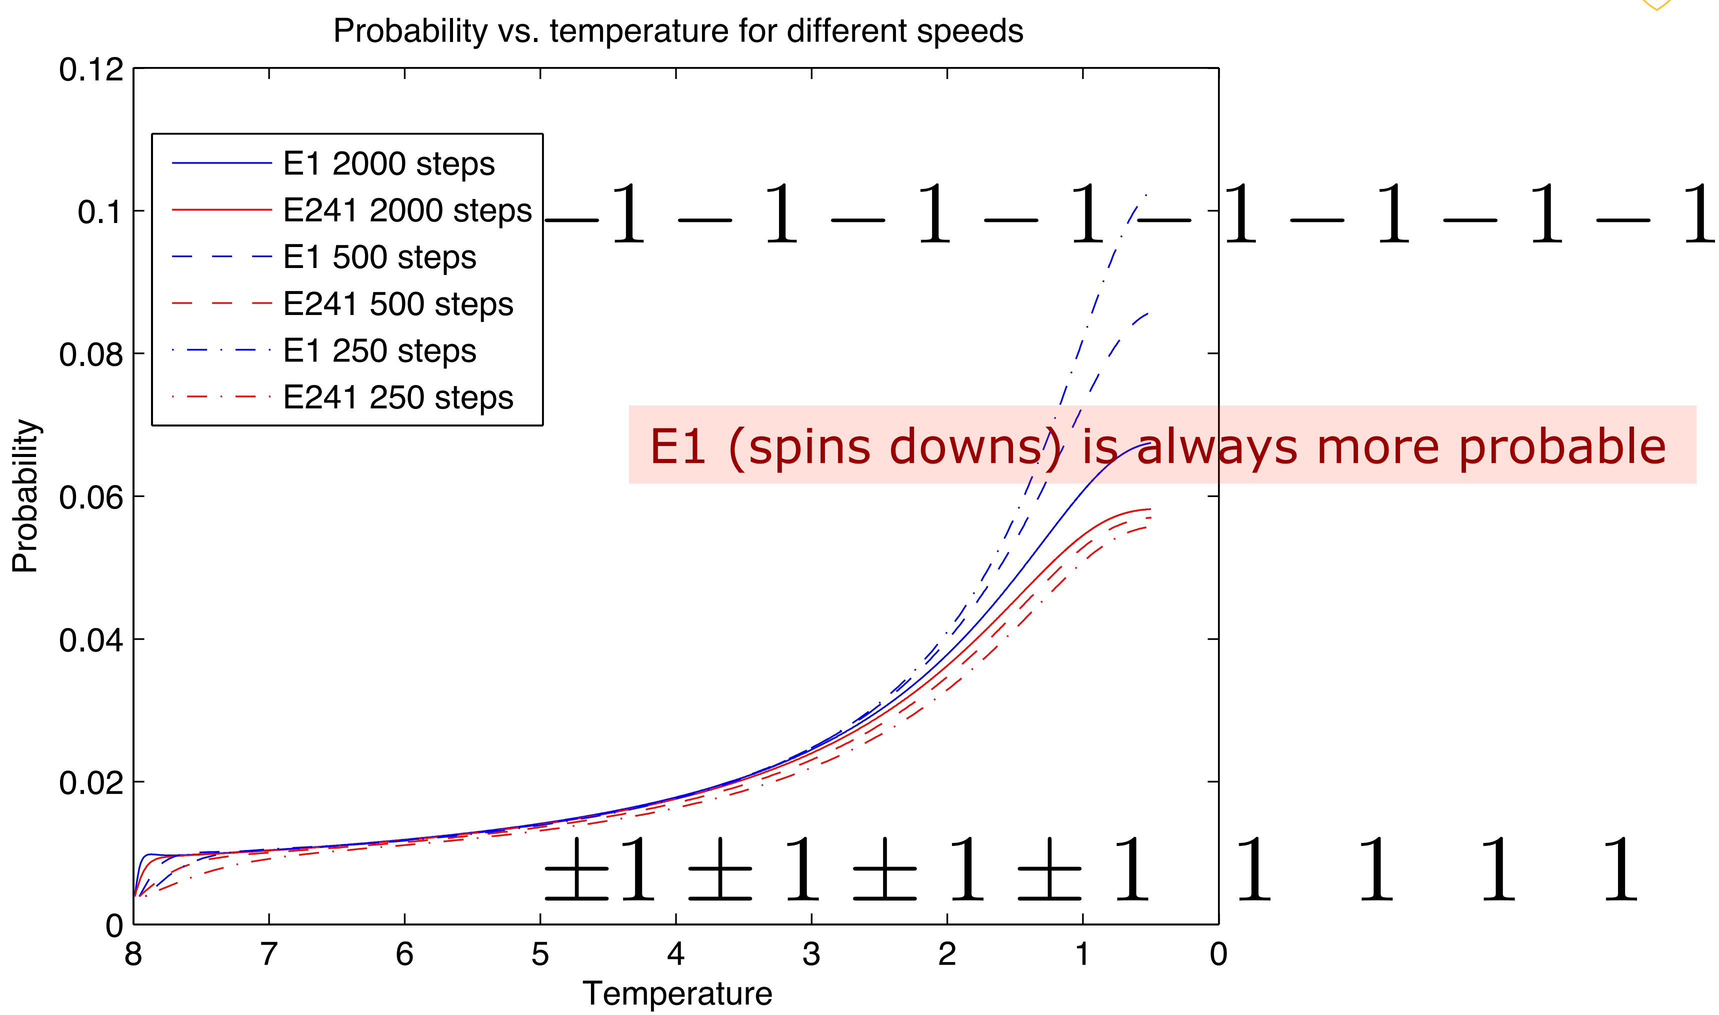

E1 (spins downs) is always more probable

Simulated Quantum Annealing

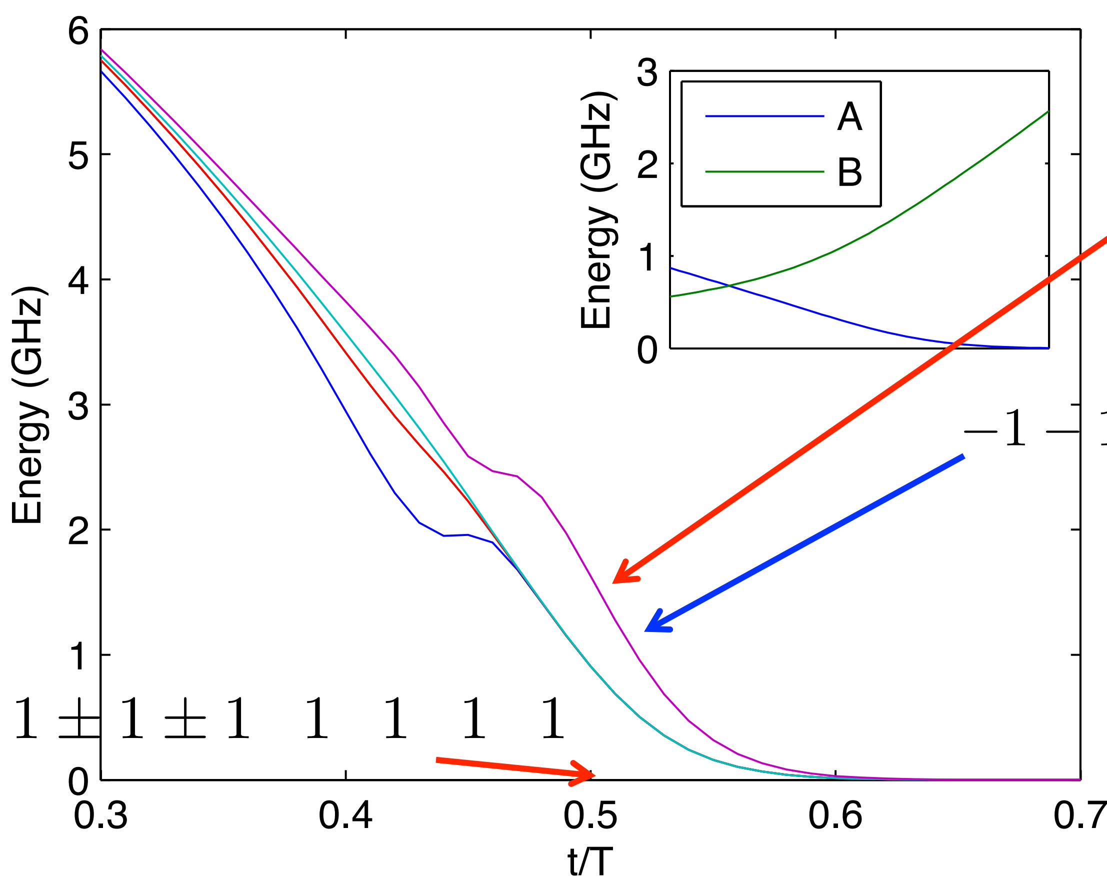

Gap 1.5 GHz

(Temp: 0.35 GHz)

$-1 -1 -1 -1 -1 -1 -1 -1$

$\pm 1 \pm 1 \pm 1 \pm 1 \quad 1 \quad 1 \quad 1 \quad 1$

D-Wave One Experiments

144 embeddings

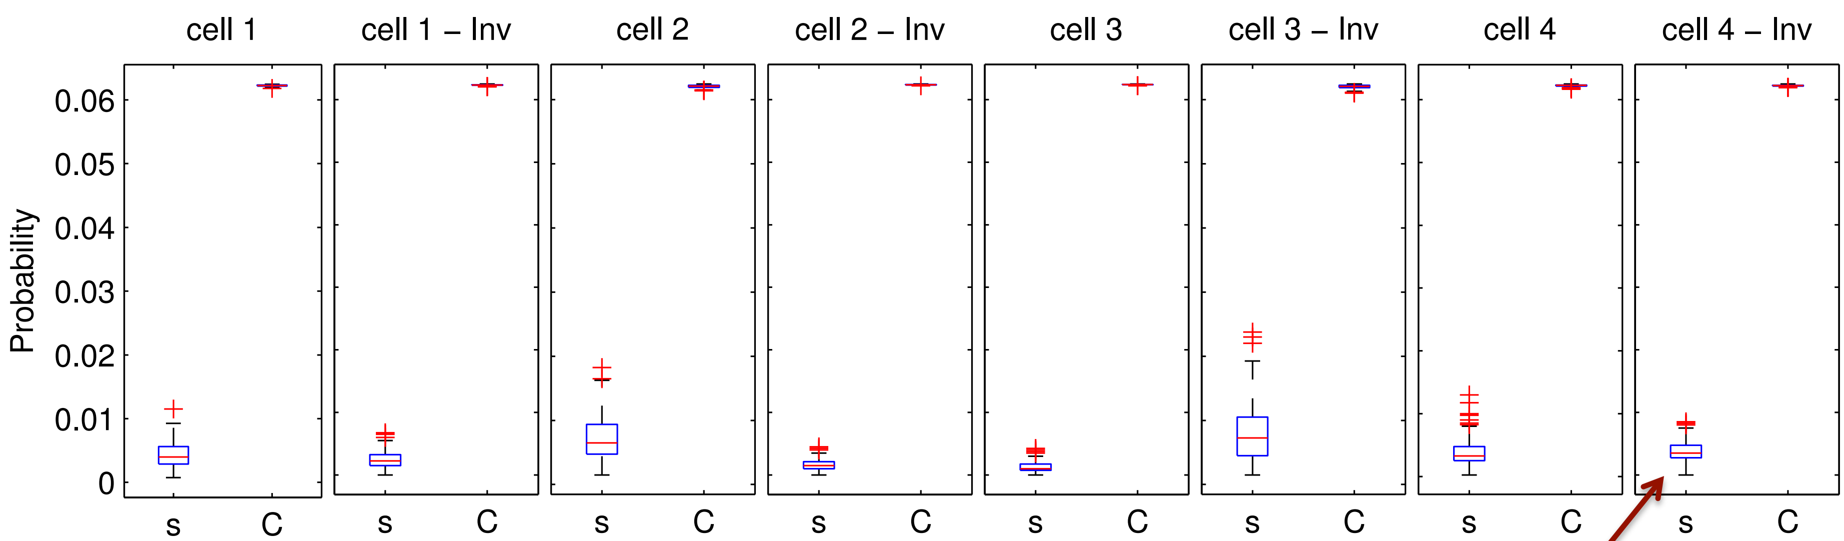

Quantum Signature: the isolated state is suppressed

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

Scalable Quantum Signature

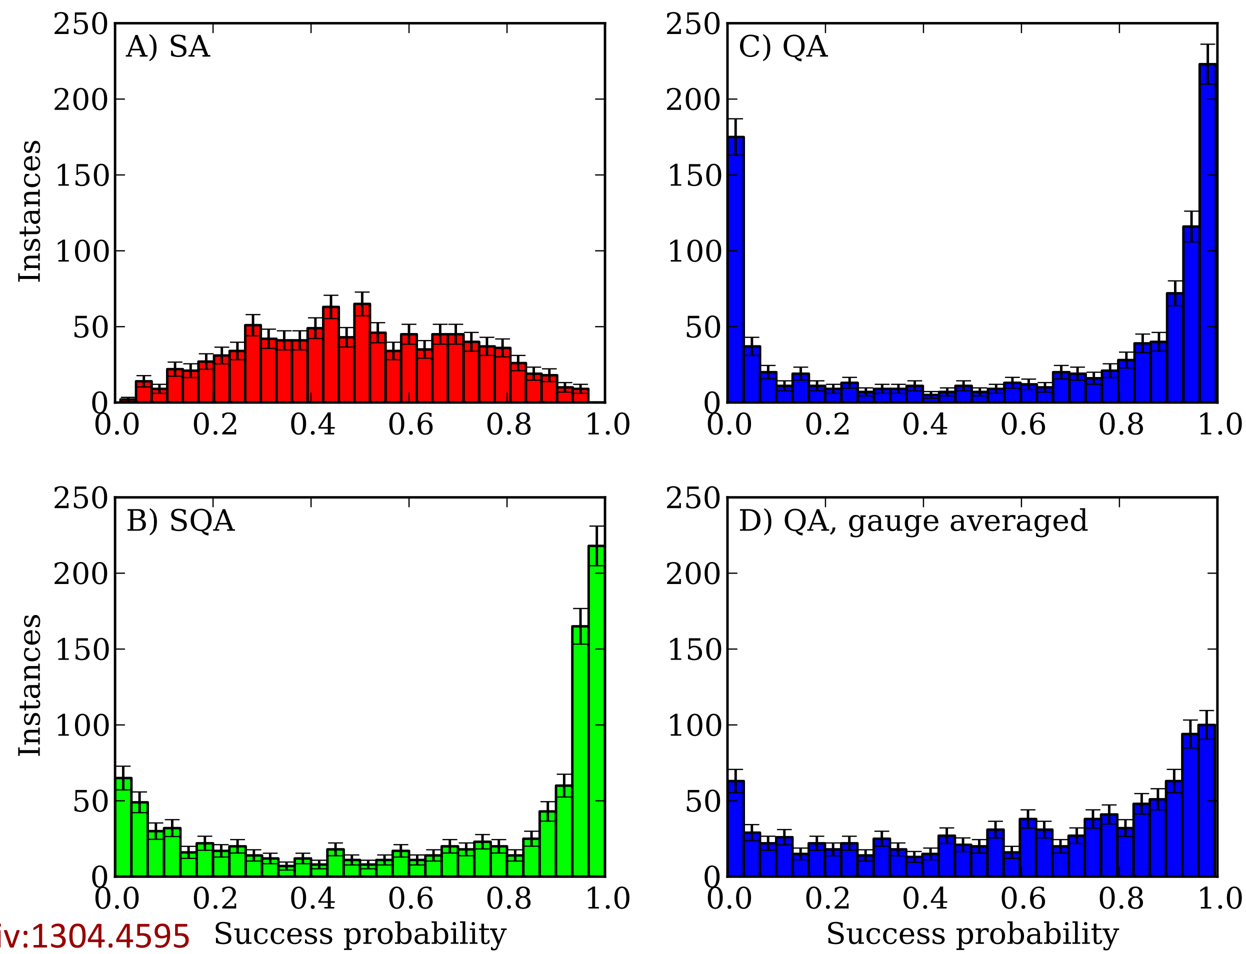

Reference: arXiv:1304.4595

Witness for Entanglement?

Collaborating with D-Wave

Only they can take the measurements needed

Currently preparing calculations for 8-qubit witness

It will take a lot of computation, O(10M) CPU hours

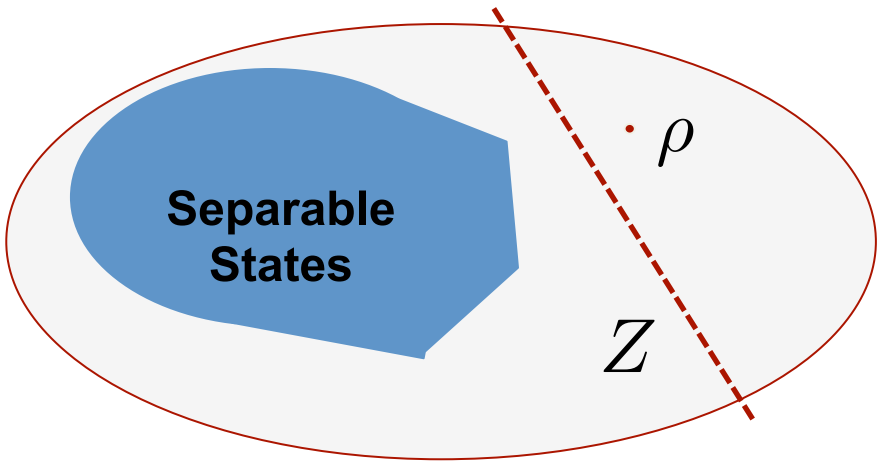

Spedaliari, AQC 2013

Performance of Classical Exact Solvers

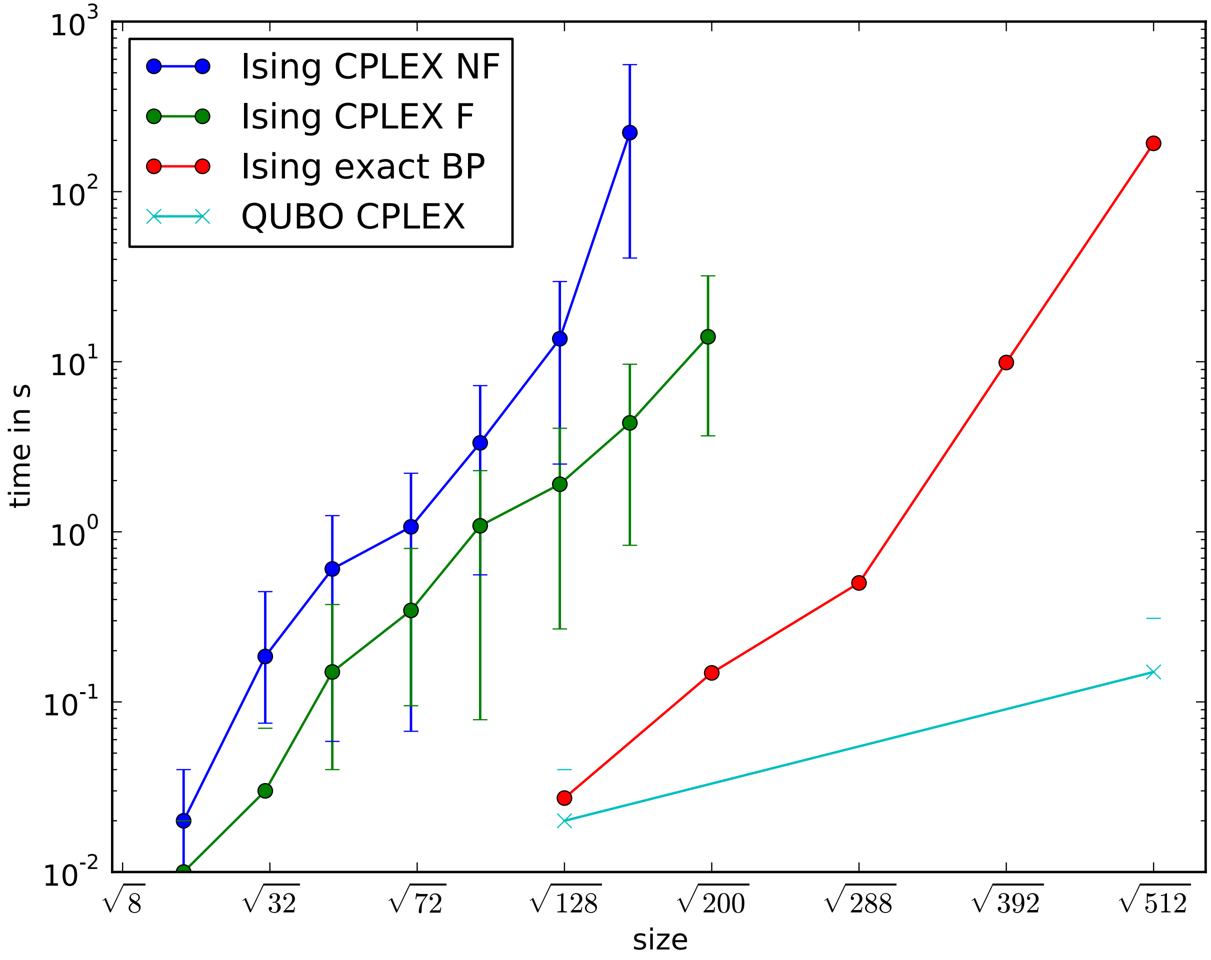

Quantum Annealing vs. Simulated Annealing

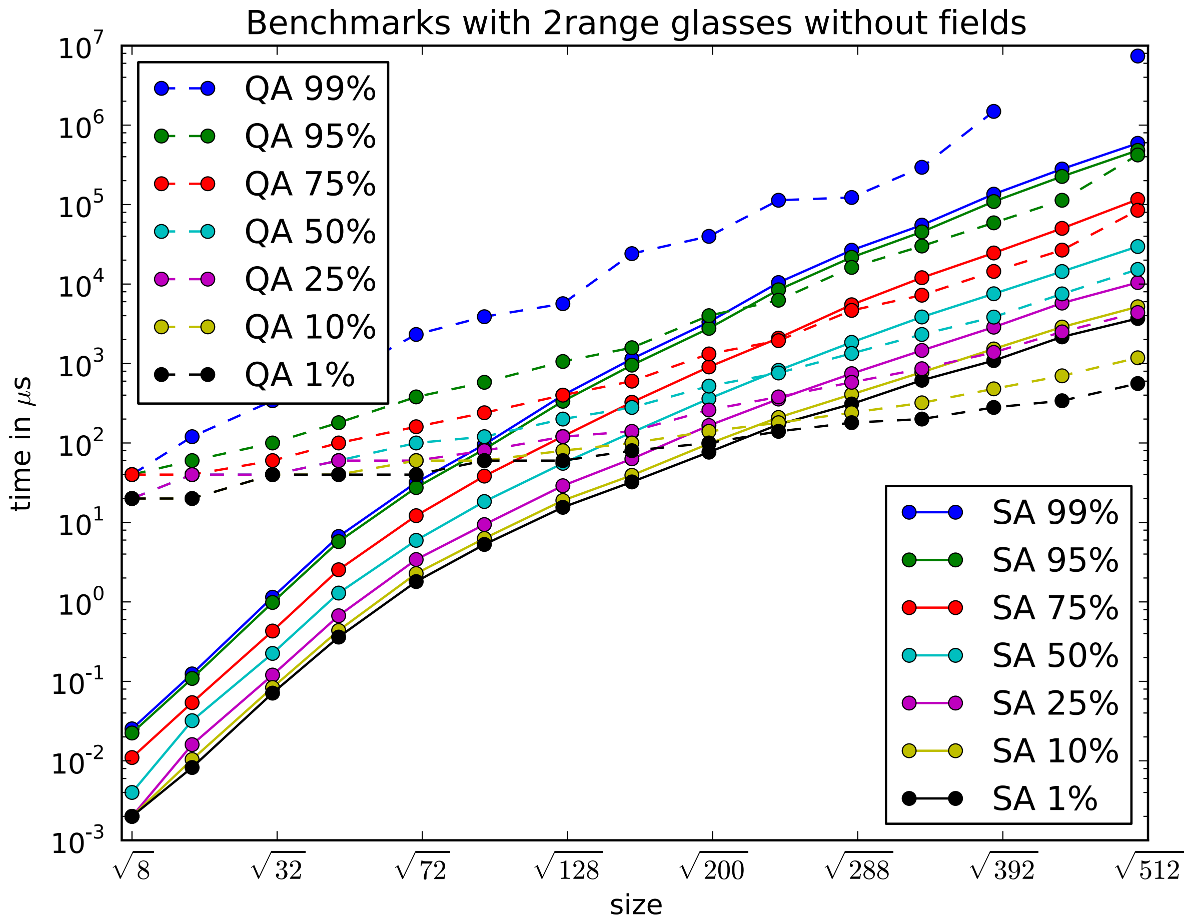

D-Wave 2 vs. Eight-core Pentium (USC & ETH)

Quantum Annealing vs. Simulated Annealing

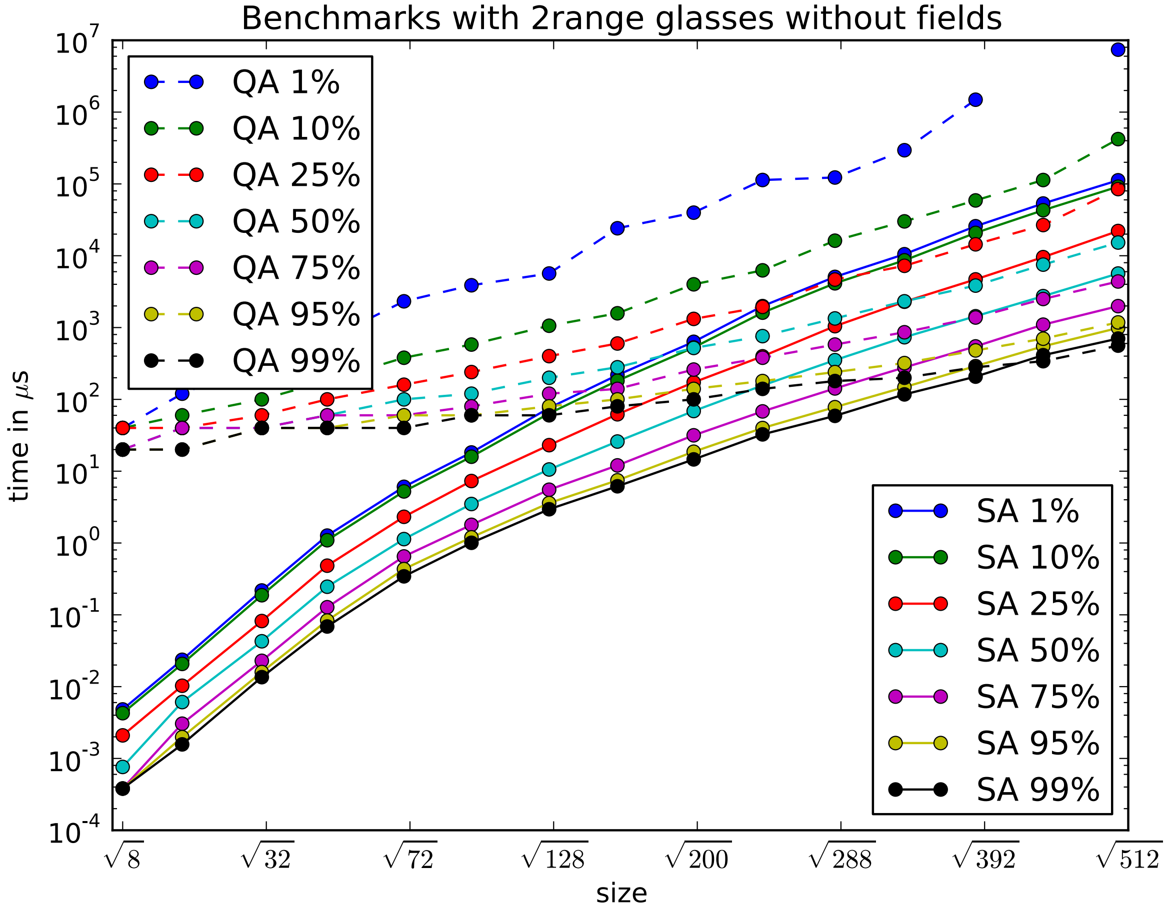

D-Wave 2 vs. Nvidia Kepler GPU (USC & ETH)

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

Software Verification and Validation

Applying Machine Learning

Collaboration with 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

Sensor assignment

Tracking

Reordering is an Algorithmic Bottleneck

T i m i n g   i n f o r m a t i o n

	CPU(seconds)	%CPU	Clock(seconds)	%Clock
Initialization	4.2627E+02	1.97	5.3379E+01	1.97
Element processing ...	1.4466E+03	6.69	1.8085E+02	6.68
Binary databases	7.4405E+00	0.03	9.2917E-01	0.03
ASCII database	3.0924E-05	0.00	3.1000E-05	0.00
Contact algorithm	1.2142E+02	0.56	1.5175E+01	0.56
Interf. ID 1	2.3144E+01	0.11	2.8875E+00	0.11
Interf. ID 2	2.2887E+01	0.11	2.8631E+00	0.11
Interf. ID 3	2.3338E+01	0.11	2.9167E+00	0.11
Interf. ID 4	2.2896E+01	0.11	2.8623E+00	0.11
Interf. ID 5	2.9136E+01	0.13	3.6410E+00	0.13
Contact entities	0.0000E+00	0.00	0.0000E+00	0.00
Rigid bodies	5.4509E-04	0.00	5.5300E-04	0.00
Implicit Nonlinear ...	9.6004E+01	0.44	1.1997E+01	0.44
Implicit Lin. Alg. ...	1.9539E+04	90.30	2.4431E+03	90.30
T o t a l s	2.1637E+04	100.00	2.7055E+03	100.00

LS-DYNA, half million elements, eight threads

Reordering with Metis WND took ~43 sec., ~1.6% of total time

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

Simplex Method

Interior Point Method

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

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

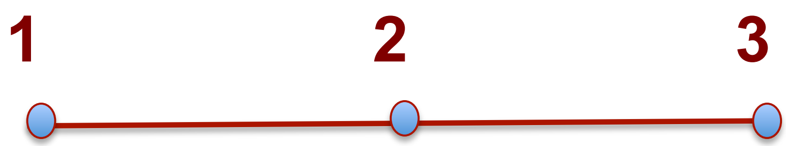

$$\underbrace{\begin{bmatrix} - & - & - \\ - & - & + \\ - & + & - \\ - & + & + \\ + & - & - \\ + & - & + \\ + & + & - \\ + & + & + \end{bmatrix}}_{\text{Solutions}^{T}} * \underbrace{\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}}_{\mathbf{H}} * \underbrace{\begin{bmatrix} - & - & - & - & + & + & + & + \\ - & - & + & + & - & - & + & + \\ - & + & - & + & - & + & - & + \end{bmatrix}}_{\text{All possible solutions}} = \underbrace{\begin{bmatrix} 5 \\ 1 \\ 5 \\ 1 \\ 1 \\ 5 \\ 1 \\ 5 \end{bmatrix}}_{\text{Energy}}$$

Direct embedding of Ising Model to Chimeral Graph

Embedded 10-Node Line

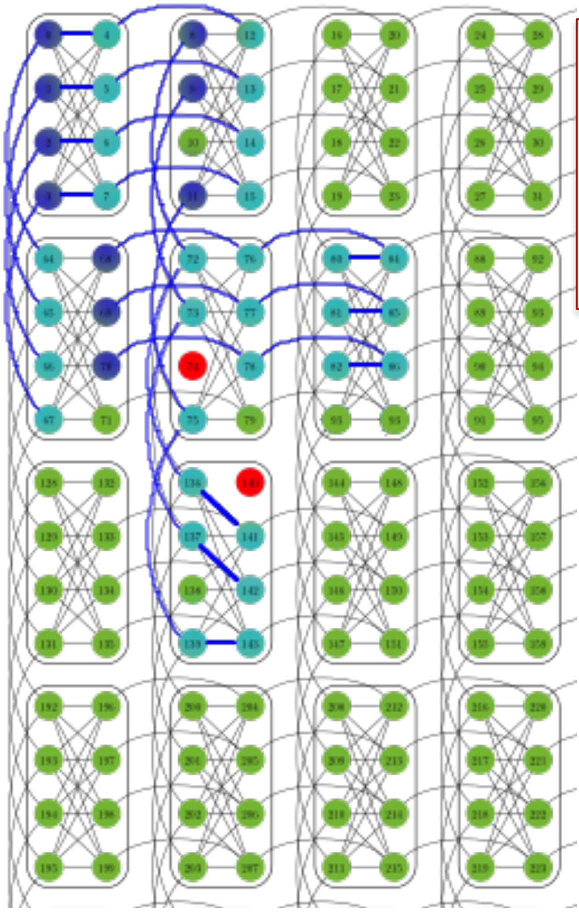

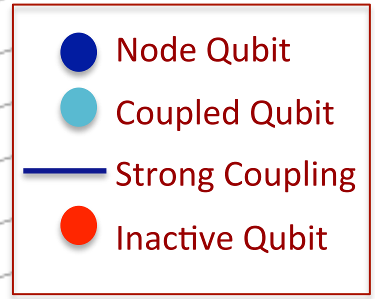

Use strong couplings between qubits to create virtual nodes with full connectivity

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

85 Node Graph Test 1: Black Box Results

H_A = Difference in Partition sizes
H_B = Number of edges crossed
$H = 2 \cdot H_A^2 + H_B$

Time(s)	Energy	H_A	H_B	H
1	177.5	1	109	111
5	179.9	1	115	117
10	164.7	1	77	79
50	150.7	1	42	44
100	147.9	1	35	37
200	143.9	1	25	27
500	140.3	1	16	18
750	142.3	1	21	23

(Note: Because the number of nodes is odd, $|H_A|$ will always be != 0)

Adjacency Matrix of 85 Node Graph

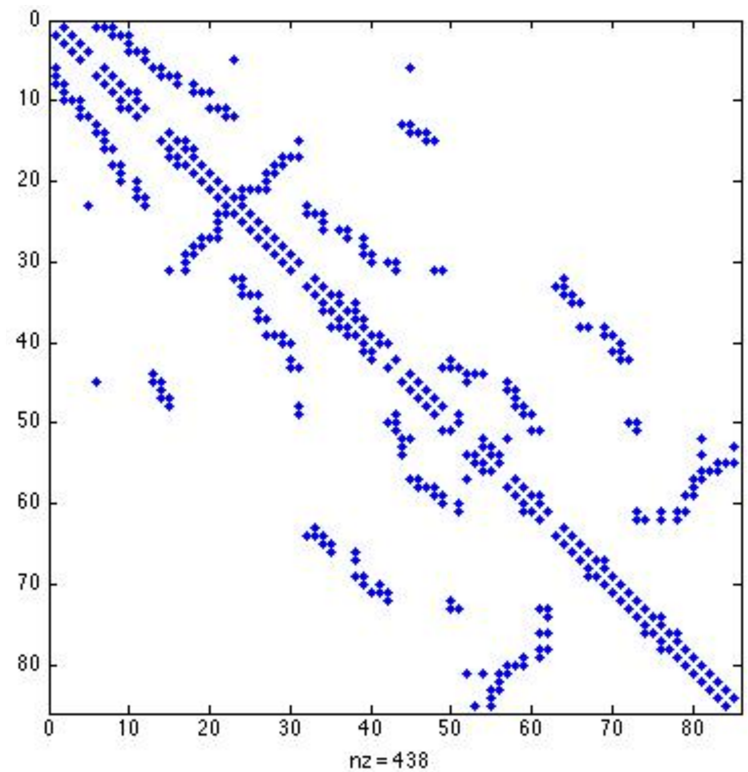

D-Wave's Black Box can produce low energy partitions of complex graphs, but at the expense of time

85 Node Graph Test 2: Sequential Embedding

Greedy embedding algorithm used to sequentially map dense Hamiltonian to D-Wave Chimeral topology

Time(s)	Post-Processing Steps	Energy	H_A	H_B	H
15.48	0	-1406.4	-1	107	109
19.56	1	-1405.2	1	115	117
21.51	2	-1413.2	1	79	81
26.62	5	-1423.2	-1	54	56
38.89	10	-1432.4	-1	31	33
46.57	25	-1432.8	-1	30	32
49.43	50	-1435.6	-1	23	25
50.42	200	-1436.8	-1	20	22

Results Produced by Black Box versus Sequential Embedding

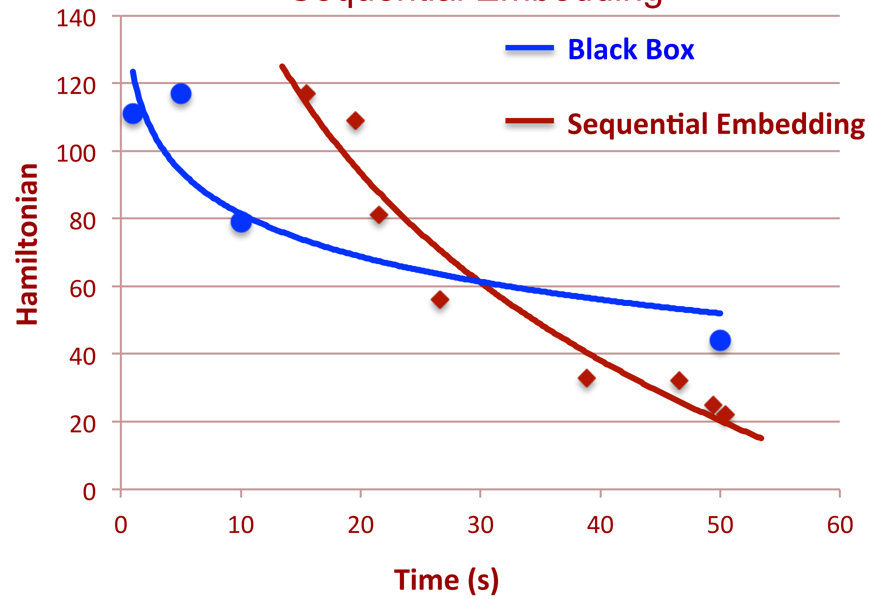

Classical Post-Processing of solutions drastically improves partitions quality for complex topologies

Sequential Embedding produces high quality results with lower time cost than Black Box

Outline

Historical Perspective

Tomorrow's Uncertainty

Adiabatic Quantum Annealing

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

Broader Open Research Problems

Quantum Optimization

New class of algorithms of which annealing is a subset

Could transform “big data” analytics

  Google’s “killer app”

Adiabatic Quantum Computing

D-Wave is an annealer, not a general purpose computer

What other phenomenon are there besides flux-qubits?

Circuit Model Quantum Computers

Goal of most of the quantum computing research community

We are entering a period of great uncertainty about the future of High Performance Computing

This creates great challenges and opportunities for DOE computational scientists

You will have a chance to shape this future

New computing systems
New mathematical algorithms
Ultimately, unimaginable scientific discoveries

Back Up

Quantum Annealing vs. Classical Annealing

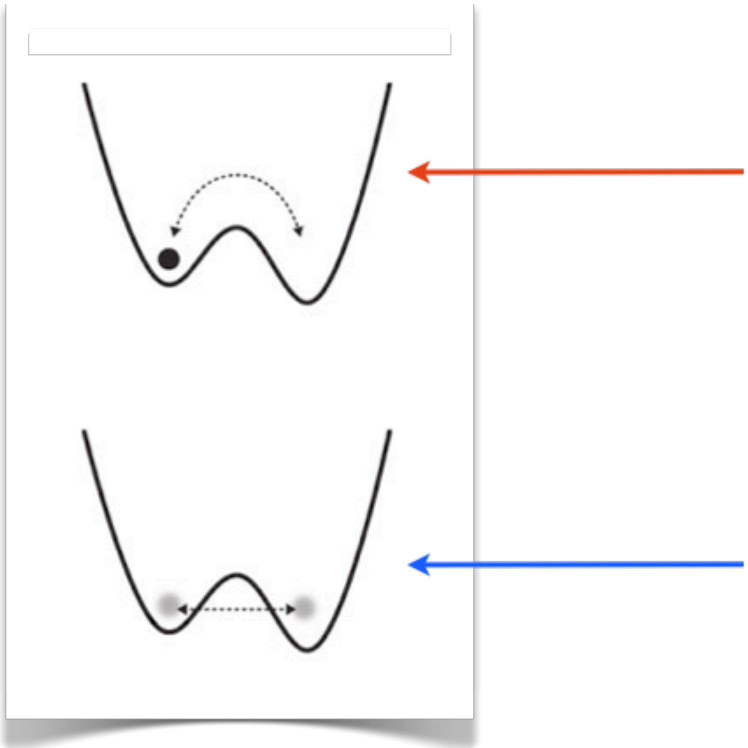

Classical annealing:
The state must jump over the barrier using temperature

Quantum annealing:
The state can go through the barrier using quantum tunneling

Recently confirmed for D-Wave´s processor

M. W. Johnson et al., "Quantum annealing with manufactured spins," Nature, vol. 473, no. 7346, pp. 194-198, May. 2011.

Theory Predicts Quantum Speed-Up

- A general version of quantum annealing is universal quantum computation (Shor's algorithm).

- We have shown that the general version gives a quadratic speed-up of simulated annealing

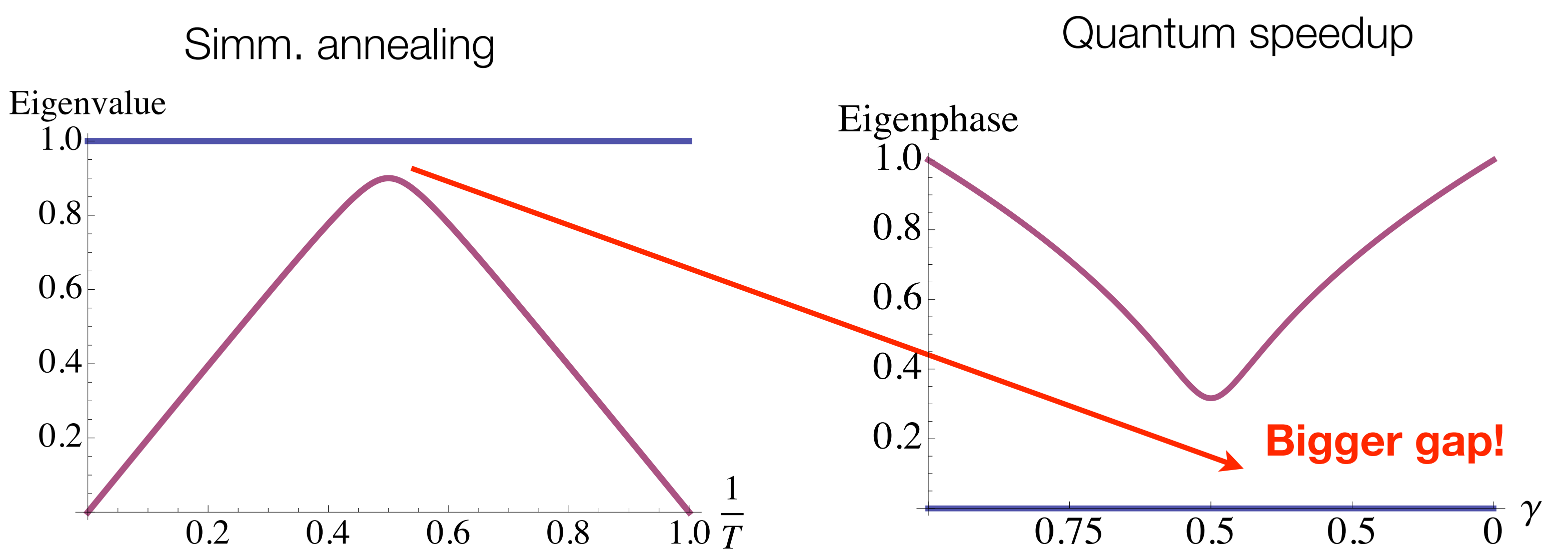

R. D. Somma, **SB**, and H. Barnum, arXiv:0712.1008, 2007.
R. D. Somma, **SB**, H. Barnum, and E. Knill, PRL, 2008.

Error Correction

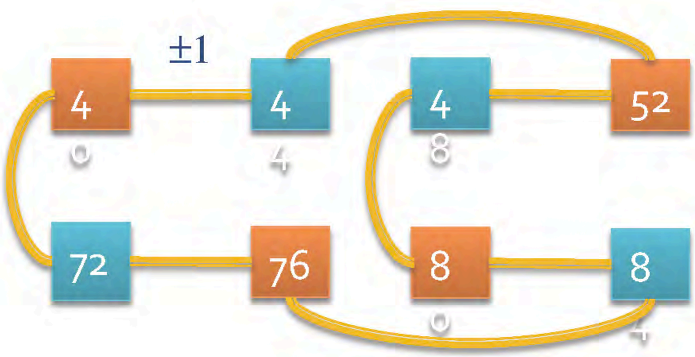

unencoded embedding

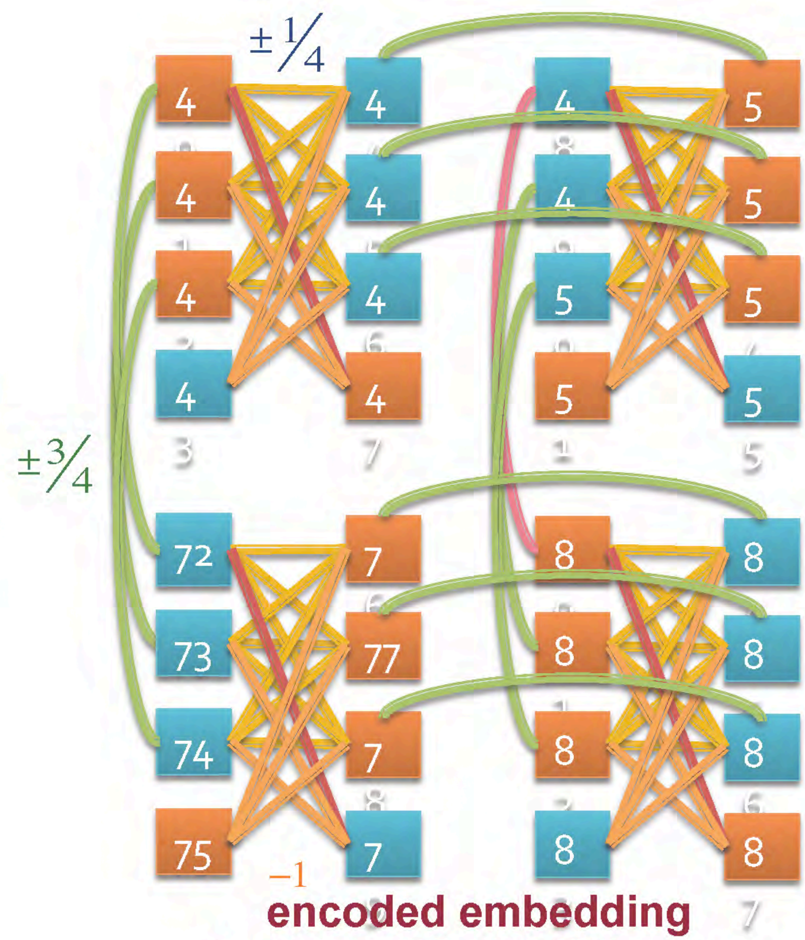

encoded embedding
bottom row ancilla qubits

In preparation

Ferromagnetic Chain Experimental Results

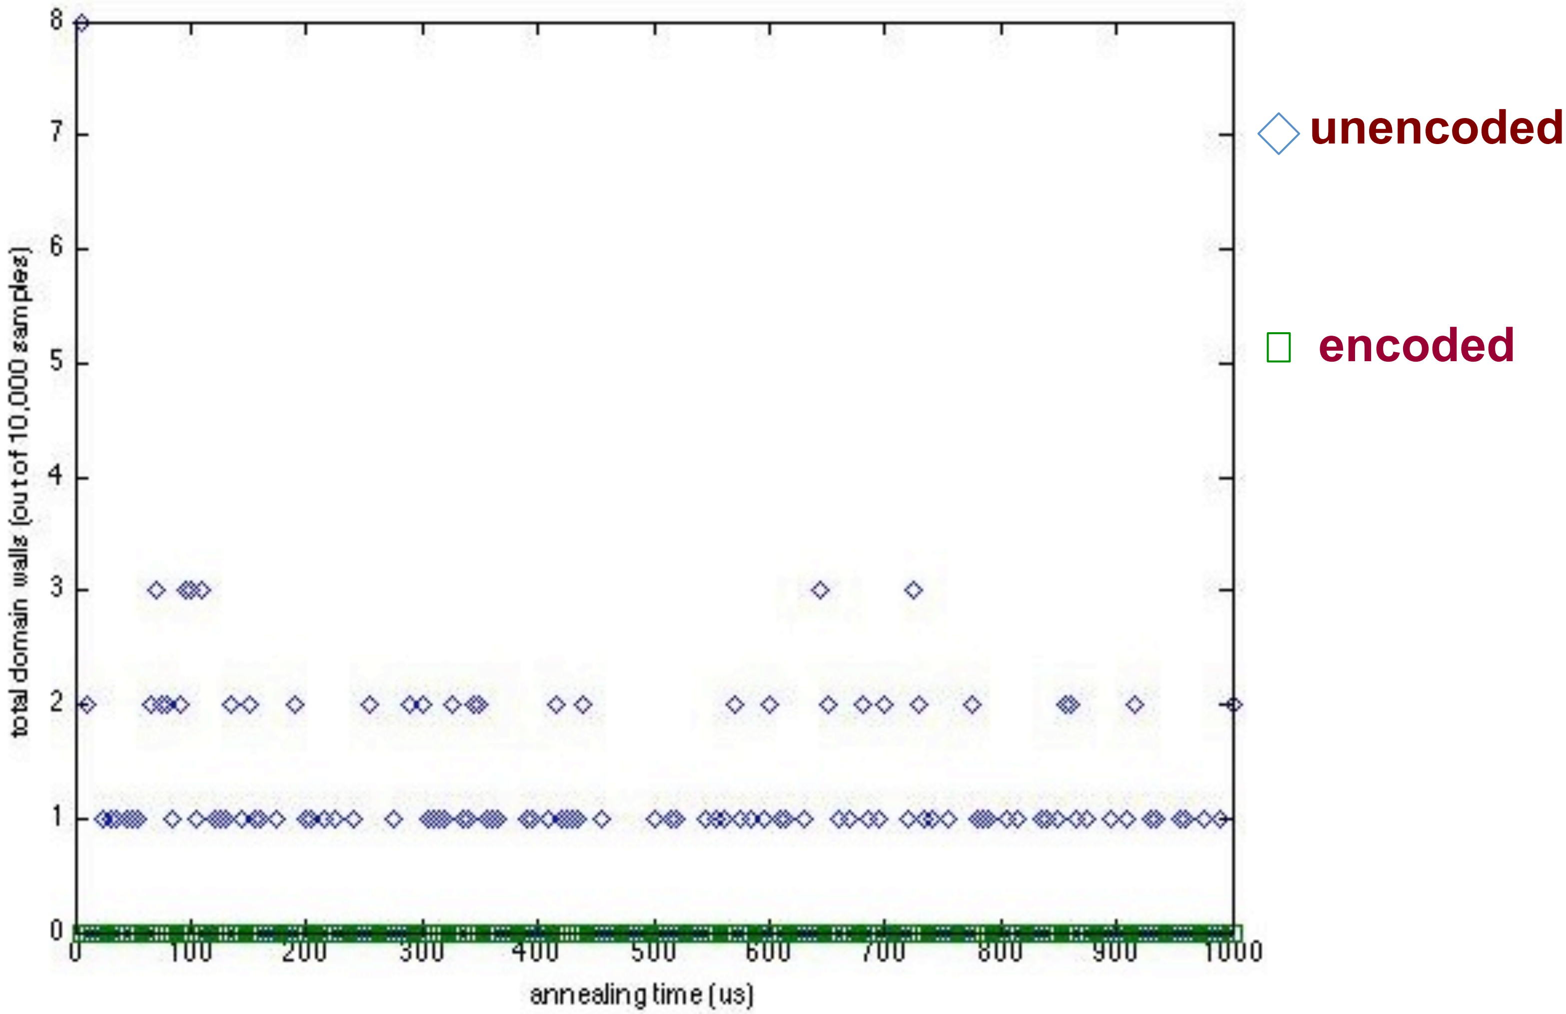

Case Study: 4-city Traveling Salesman (LM 2011)

- 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

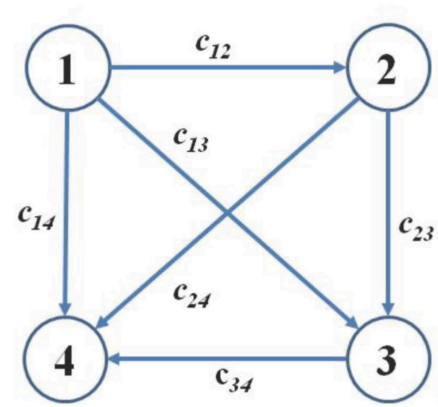

Link Costs

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

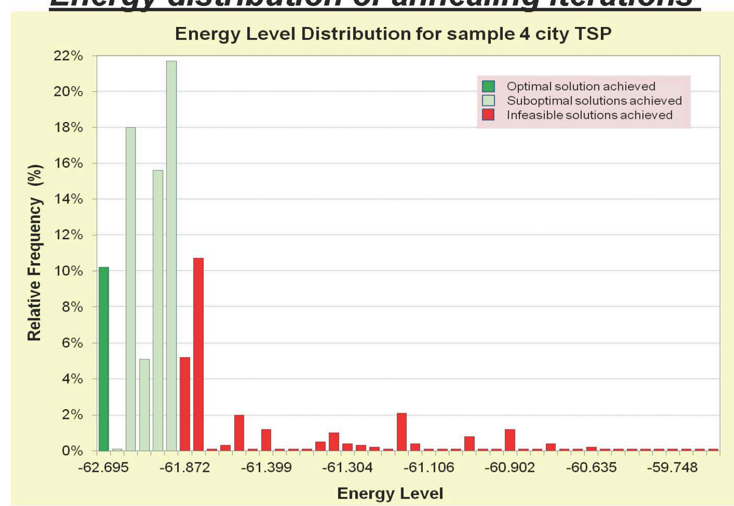

Slide Courtesy of Steve Adachi, Lockheed Martin

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

Amino-acid sequence	Interaction	ΔE
P–S–V–K–M–A	P ≡ K	-1
	P ≡ A	-2
	S ≡ M	-3
	V ≡ A	-4

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

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

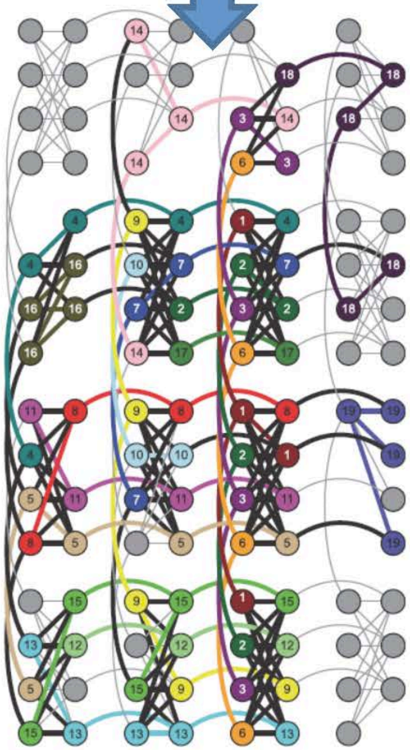

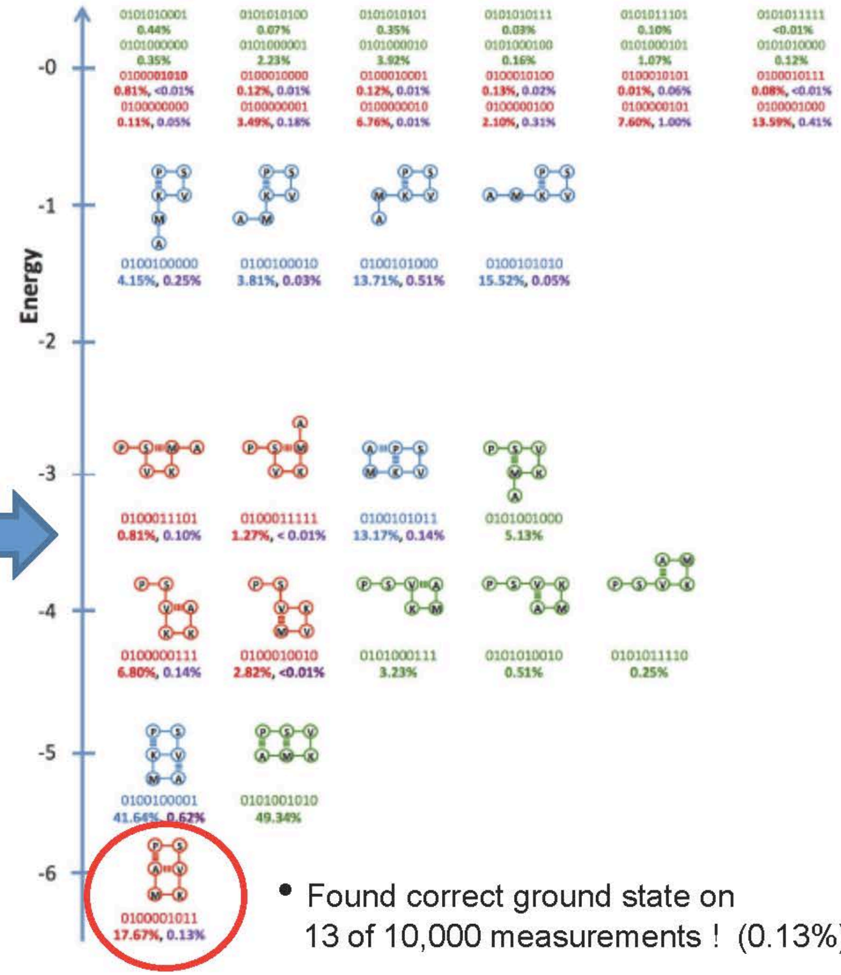

- Found correct ground state on 13 of 10,000 measurements ! (0.13%)

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