

Numerical Studies of the Quantum Adiabatic Algorithm

A.P. Young



SANTA CRUZ



Work supported by



National Science Foundation
WHERE DISCOVERIES BEGIN

Talk at CCP2014, August 13, 2014, Boston University

Collaborators: I. Hen, M. Wittmann, E. Farhi, P. Shor, D. Gosset, A. Sandvik, V. Smelyanskiy, S. Knysh, M. Guidetti. Computer support from 

Why is quantum different?

- **Quantum Parallelism**

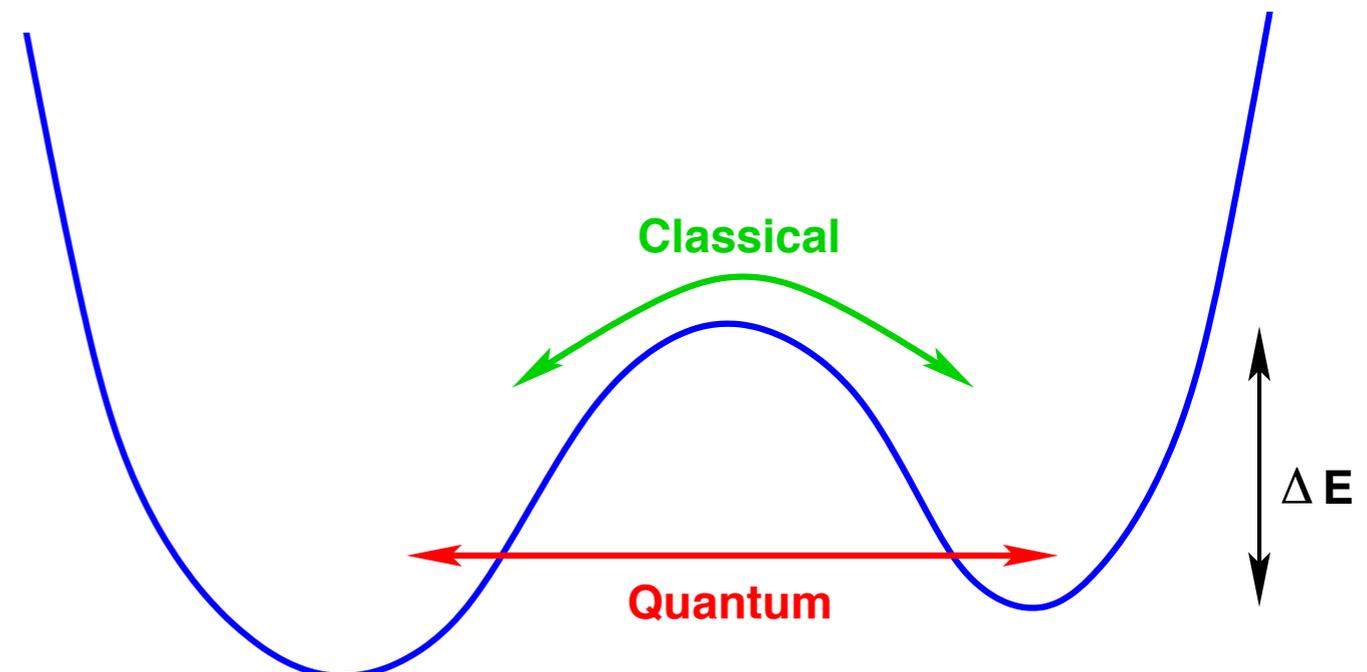
A quantum state is a linear combination of basis states. Acting with a discrete unitary transformation acts **in parallel** on all states. This is the basis of the **traditional paradigm of quantum computing**, e.g. **Shor's algorithm** for factoring integers:

Exponential time (classical) → polynomial time (QM)

Exponential speedup!

- **Quantum Tunneling**

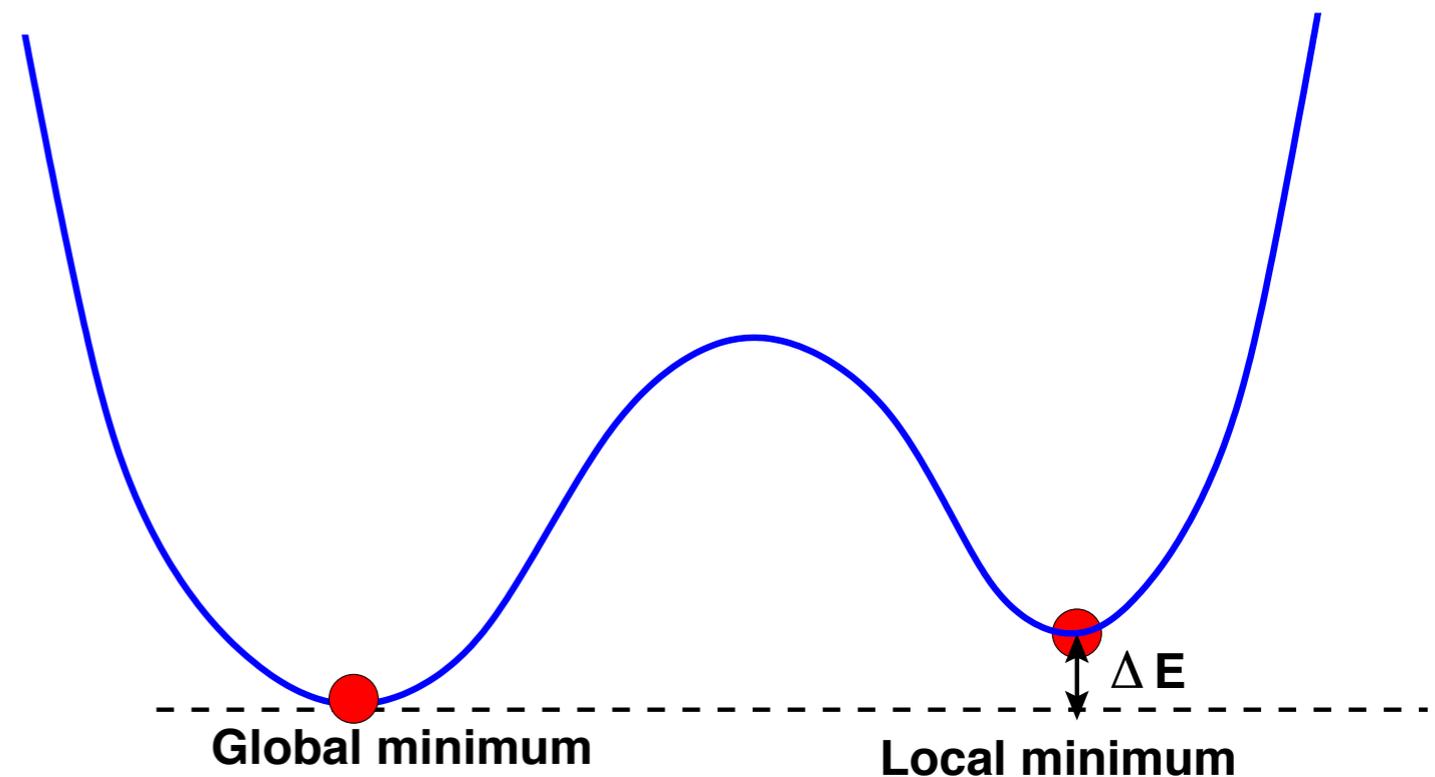
Classically, a particle has to be thermally activated over a barrier, in QM it can tunnel through it. This gives a **2nd paradigm for QC** known as ...



Quantum Annealing / Quantum Adiabatic Algorithm (for solving Optimization Problems)

Suppose we have to find the minimum of a function (Hamiltonian) of many variables, an **optimization problem**. In problems of interest no set of variables simultaneously minimizes each term. This is called “**frustration**” in the spin glass language (previous talk).

Optimization problems occur in physics, e.g. **spin glasses**, biology, e.g. **protein folding**, computer science, e.g. **SAT** problems, engineering, e.g. **pattern recognition** (interest of **Google**).



Minimizing w.r.t. each variable in turn (**greedy algorithm**) you get stuck in a **local minimum**. How can we get out of a local minimum to find the **global minimum**?

Physics inspired methods for Optimization Problems

- **Thermal (classical) annealing** (Kirkpatrick et al.)
also known as simulated annealing (**SA**). Put in a **non-zero temperature** and do (classical) Monte Carlo. Some probability to go “uphill” in energy, so can escape from local minima by **thermal activation OVER** the barriers. $T \rightarrow 0 \text{ as } t \rightarrow \infty$
- **Quantum Annealing (QA)** (Santoro et al, Kadawoki & Nishimori)
Same general idea as SA but induce **quantum** (rather than thermal) **fluctuations** by adding a simple but **non-commuting driver Hamiltonian H_D** . System then escape from local minima by **quantum tunneling THROUGH** barriers. $H_D \rightarrow 0 \text{ as } t \rightarrow \infty$
Often simulated on a classical computer using QMC. Then called **Simulated Quantum Annealing (SQA)** (Troyer et al.)
- **Quantum Adiabatic Algorithm (QAA)** (Farhi et al.)
Similar to QA but (a) always stay in the ground state of the Hamiltonian (**adiabatic theorem**) and (b) runs on a **quantum computer / quantum annealer**. (See rest of talk.)

Quantum Adiabatic Algorithm

Want to find the ground state of a **problem Hamiltonian** \mathcal{H}_P involving Ising spins, $\sigma_i^z = \pm 1$, or equivalently, bits $b_i = 0$ or 1 , e.g. a spin glass (previous talk, Helmut Katzgraber)

$$\mathcal{H}_P = - \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z - \sum_i h_i \sigma_i^z$$

Make quantum by adding a **non-commuting driver Hamiltonian**.

Simplest is a transverse field: $\mathcal{H}_D = -\Gamma \sum_i \sigma_i^x$

Total Hamiltonian:

$$\mathcal{H}(t) = [1 - s(t)]\mathcal{H}_D + s(t)\mathcal{H}_P$$

with $s(0) = 0$, $s(\tau) = 1$.



System starts in ground state of driver Hamiltonian. If process is **adiabatic** (and $T \rightarrow 0$), it ends in **g.s. of problem Hamiltonian**

Quantum Adiabatic Algorithm

Want to find the ground state of a **problem Hamiltonian** \mathcal{H}_P involving Ising spins, $\sigma_i^z = \pm 1$, or equivalently, bits $b_i = 0$ or 1 , e.g. a spin glass (previous talk, Helmut Katzgraber)

$$\mathcal{H}_P = - \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z - \sum_i h_i \sigma_i^z$$

Make quantum by adding a **non-commuting driver Hamiltonian**.

Simplest is a transverse field: $\mathcal{H}_D = -\Gamma \sum_i \sigma_i^x$

Total Hamiltonian:

$$\mathcal{H}(t) = [1 - s(t)]\mathcal{H}_D + s(t)\mathcal{H}_P$$

with $s(0) = 0$, $s(\tau) = 1$.



System starts in ground state of driver Hamiltonian. If process is **adiabatic** (and $T \rightarrow 0$), it ends in **g.s. of problem Hamiltonian**

Real Experiments:

D-Wave: large number of superconducting qubits:

D-Wave 1, 128 qubits

D-Wave 2, 512 qubits

Runs the QAA.

During the run, phase coherence is not maintained, hence call this a **quantum annealer**.

Questions:

- Does noise and non-zero T destroy quantum tunneling?
- If not, is the D-Wave machine more efficient than a classical computer?

Real Experiments:

D-Wave: large number of superconducting qubits:

D-Wave 1, 128 qubits

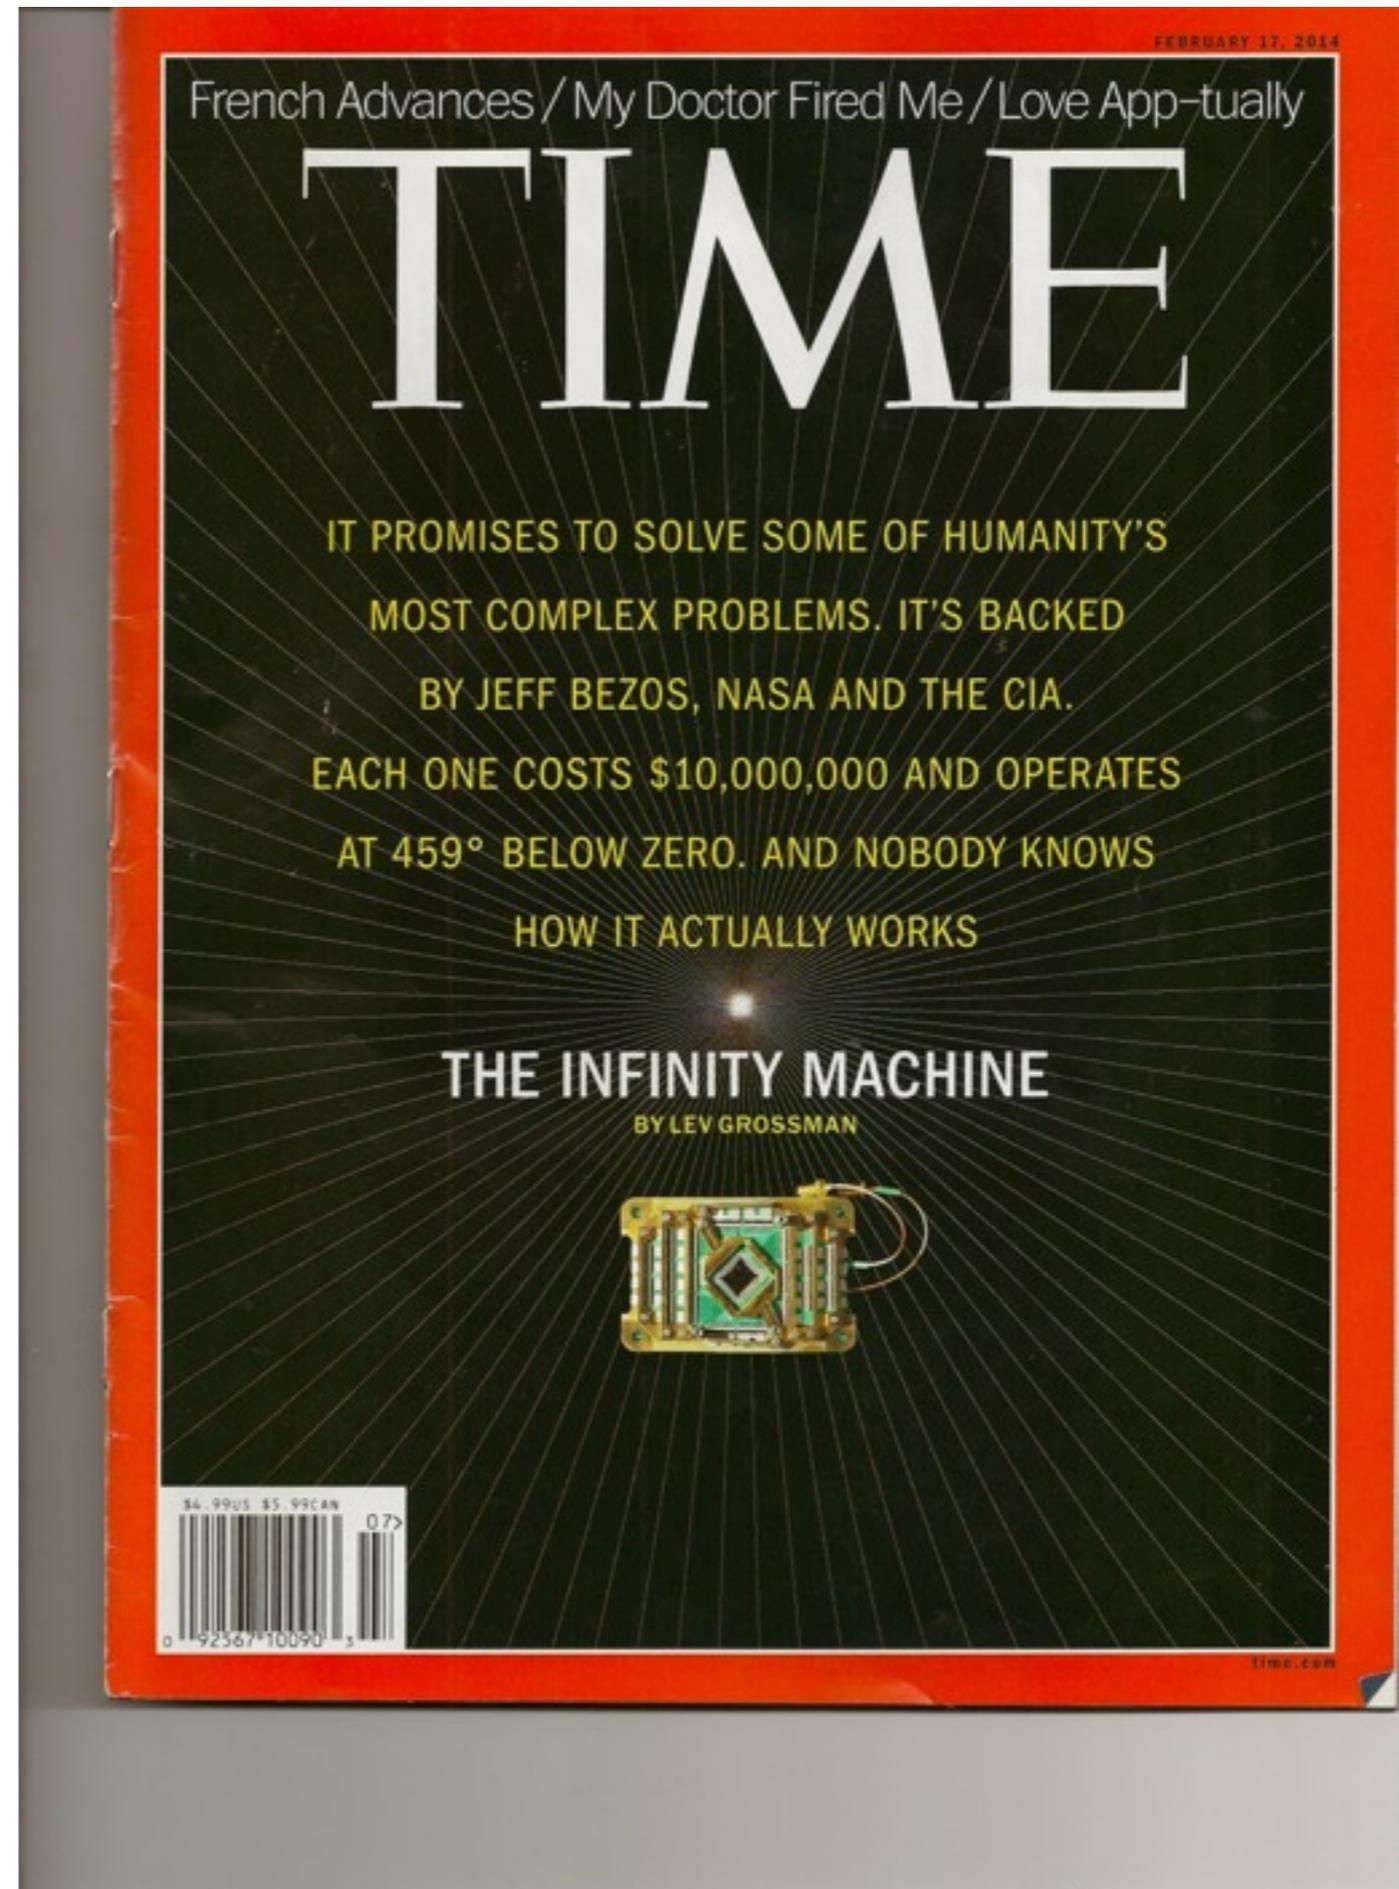
D-Wave 2, 512 qubits

Runs the QAA.

During the run, phase coherence is not maintained, hence call this a **quantum annealer**.

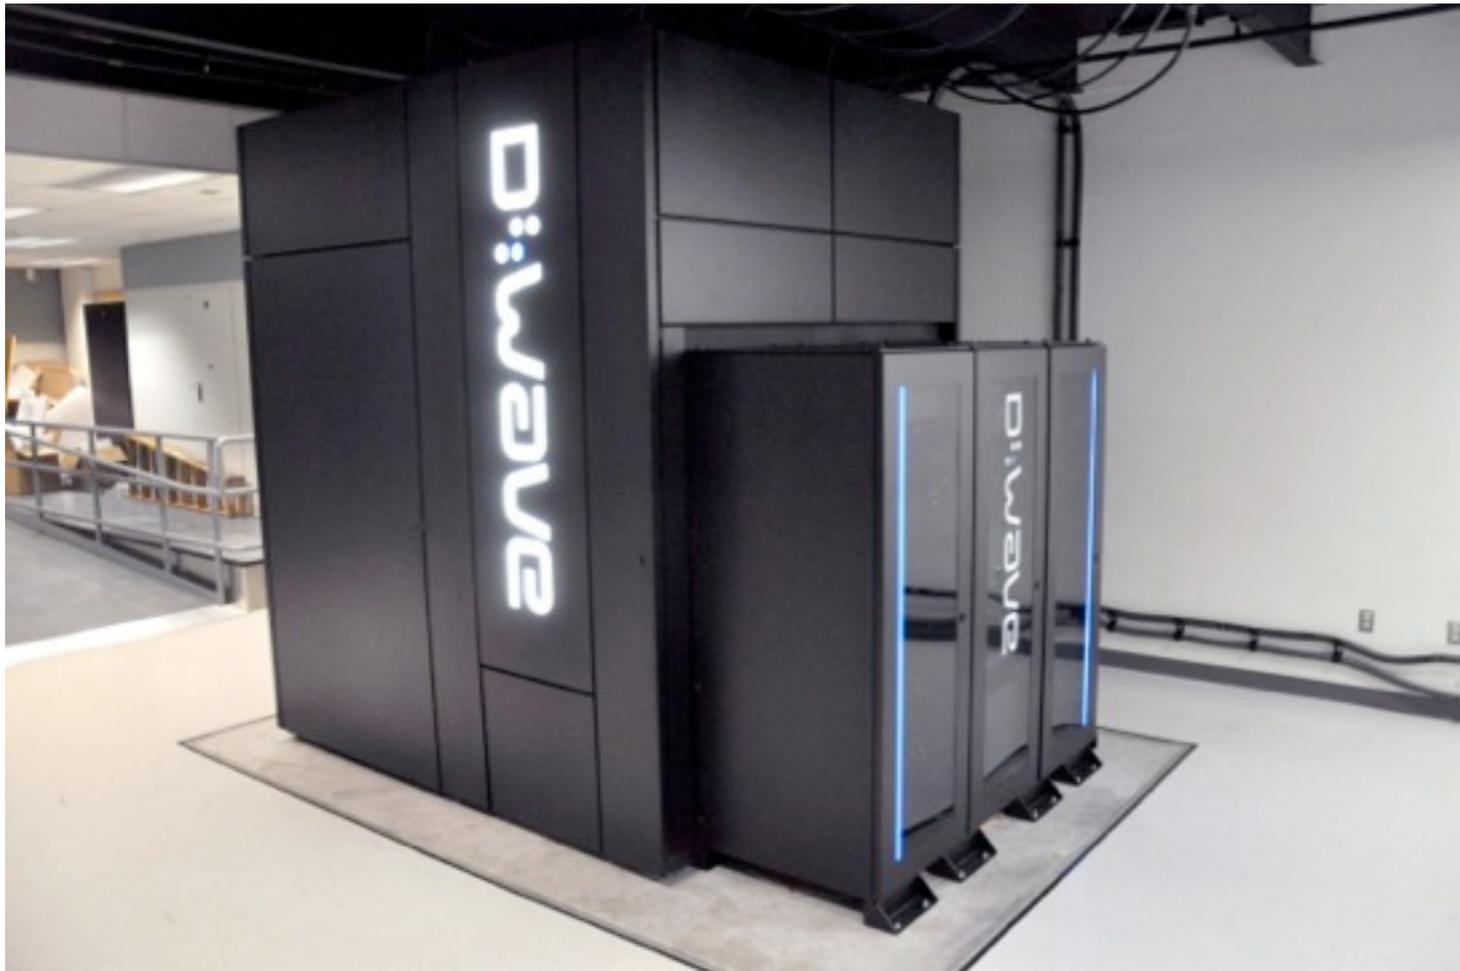
Questions:

- Does noise and non-zero T destroy quantum tunneling?
- If not, is the D-Wave machine more efficient than a classical computer?

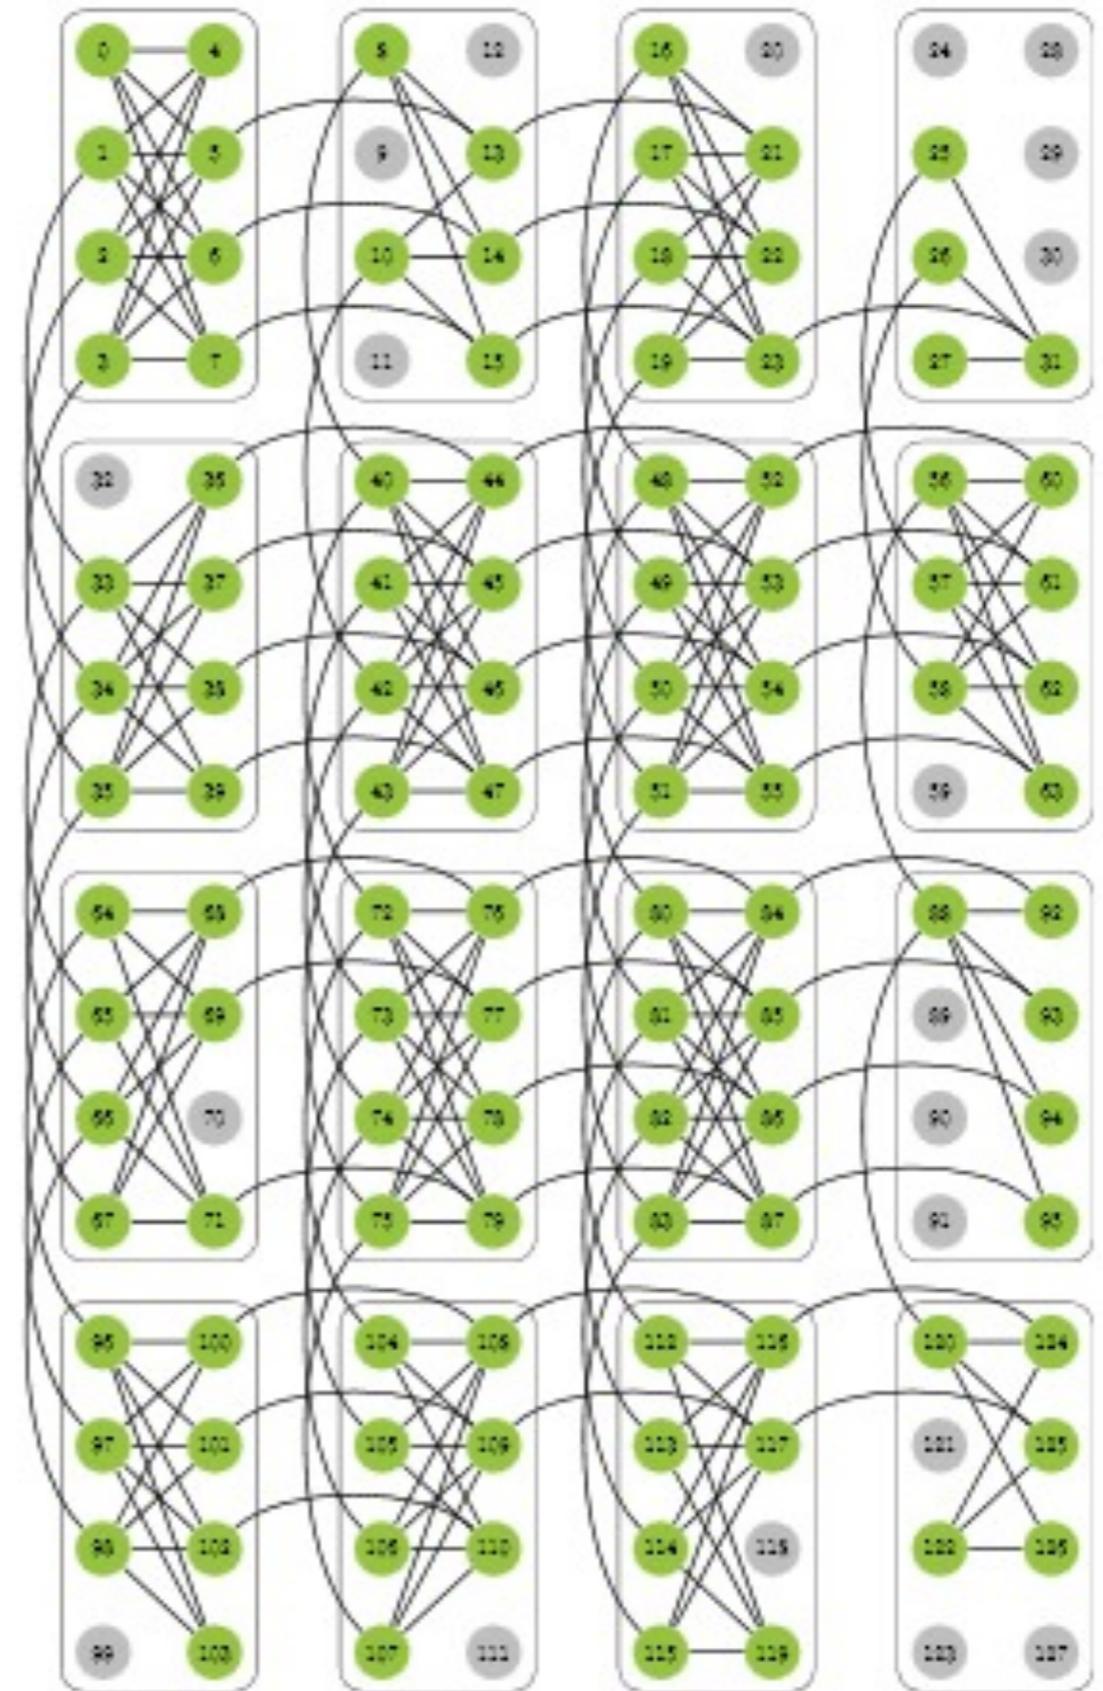
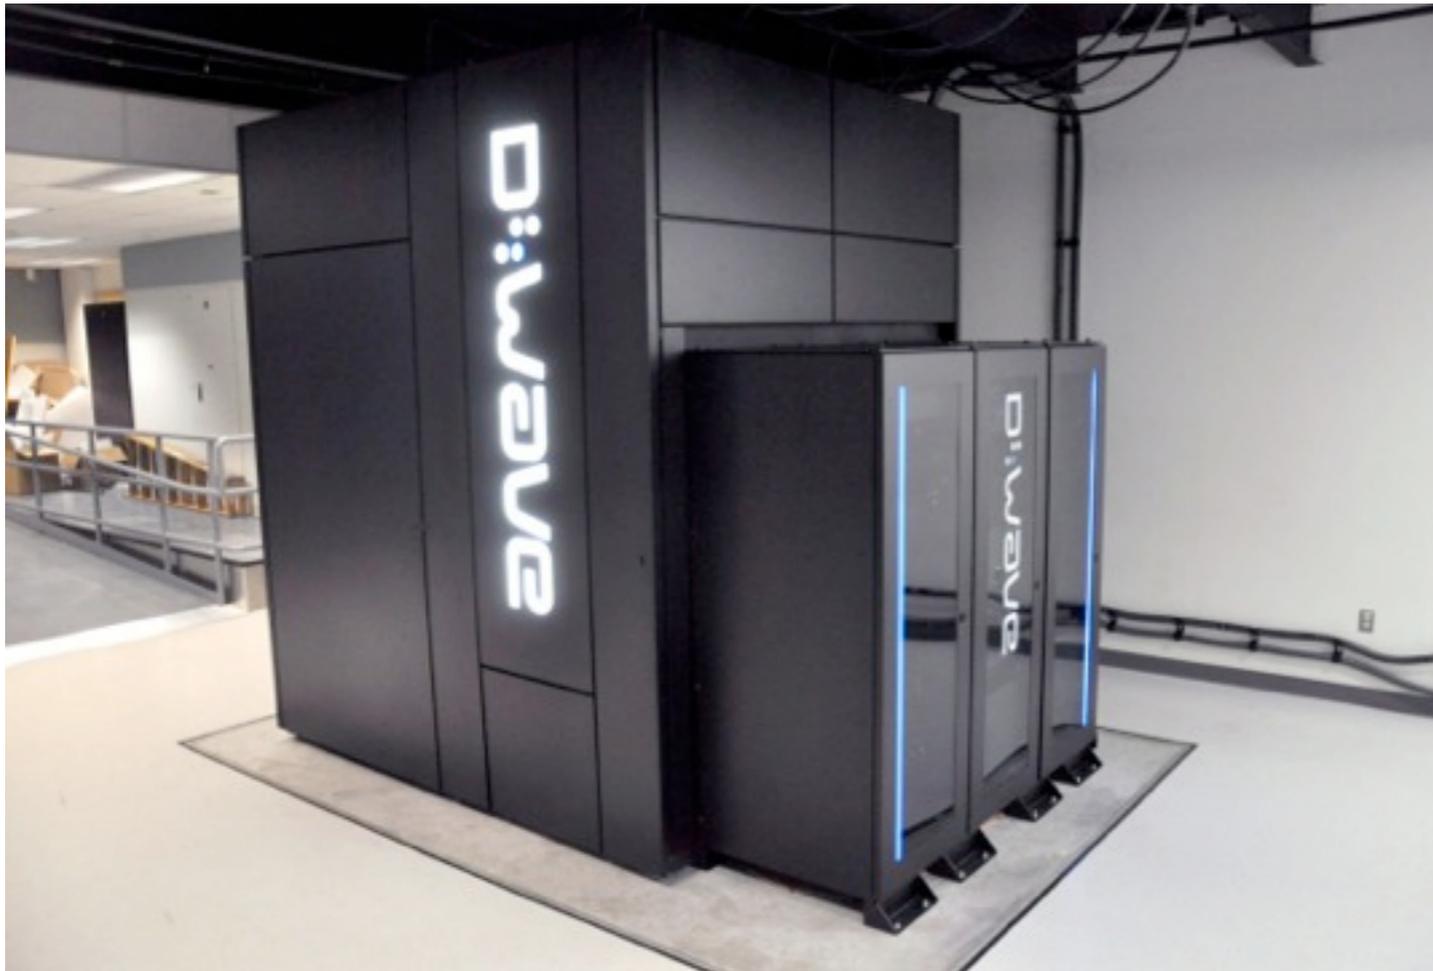


D-Wave

D-Wave



D-Wave



Connections of the qubits form a (2-dimensional) “chimera” graph, see figure for D-Wave 1.

If problem does not naturally fit on the chimera, need auxiliary qubits in which case

$$N_{qubit} > N \text{ (size of problem)}$$

Complexity

On a classical computer, the running time for hard problems increases exponentially with problem size N , i.e.

$$\tau \propto \exp(\mu N)$$

Can QA/QAA do better? Unlikely to become polynomial but can it make μ smaller?

In the traditional “gate model” paradigm of QC it is known for certain problems, e.g. integer factoring, that there is a quantum speedup, (if decoherence can be eliminated, which is a big “if”)

In the QA/QAA paradigm there is

no proof that quantum can do better, but also

no proof that it can't.

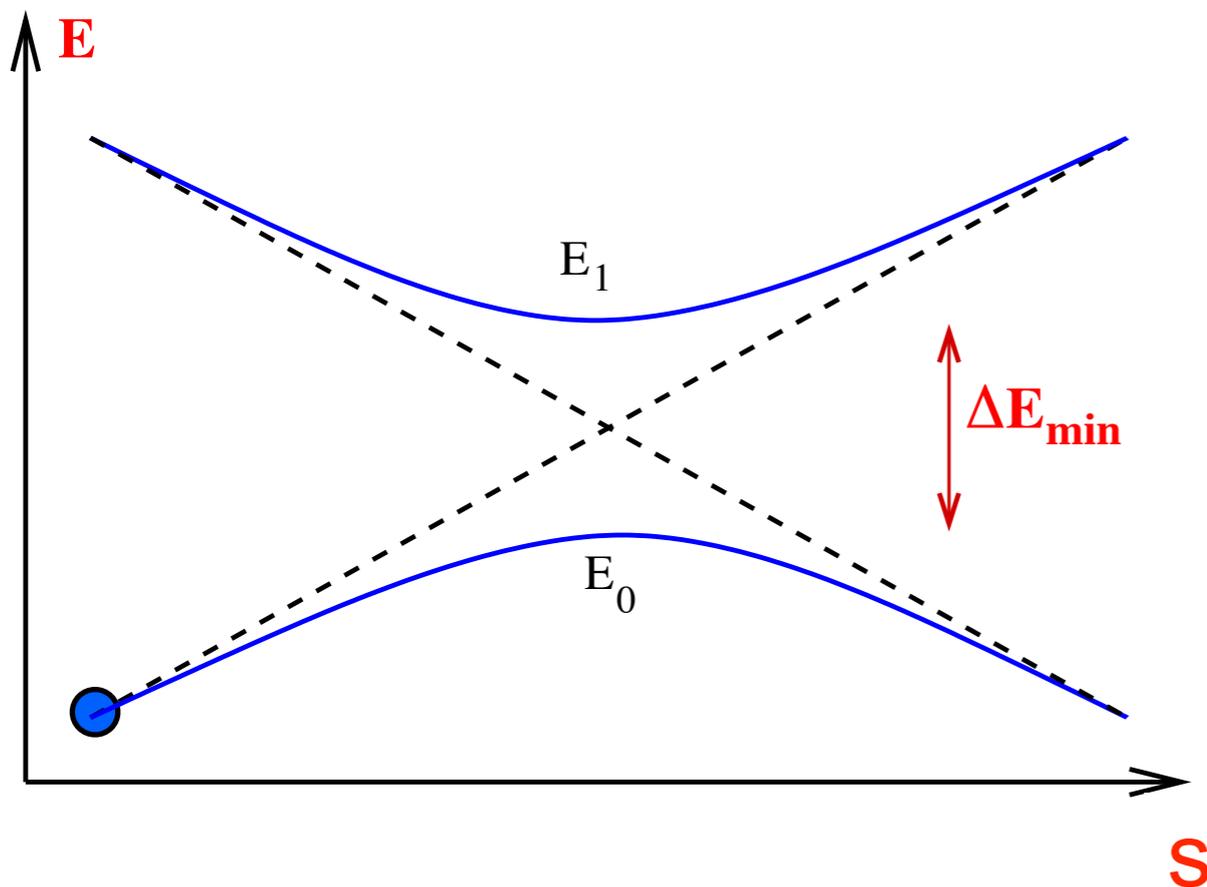
Hence we need numerics.

(There is also hope that noise/decoherence doesn't immediately remove the quantum speedup in QA/QAA (if it exists) unlike in the gate model.)

Quantum Phase Transition



Bottleneck is likely to be a **quantum phase transition (QPT)** where the **gap to the first excited state is very small**

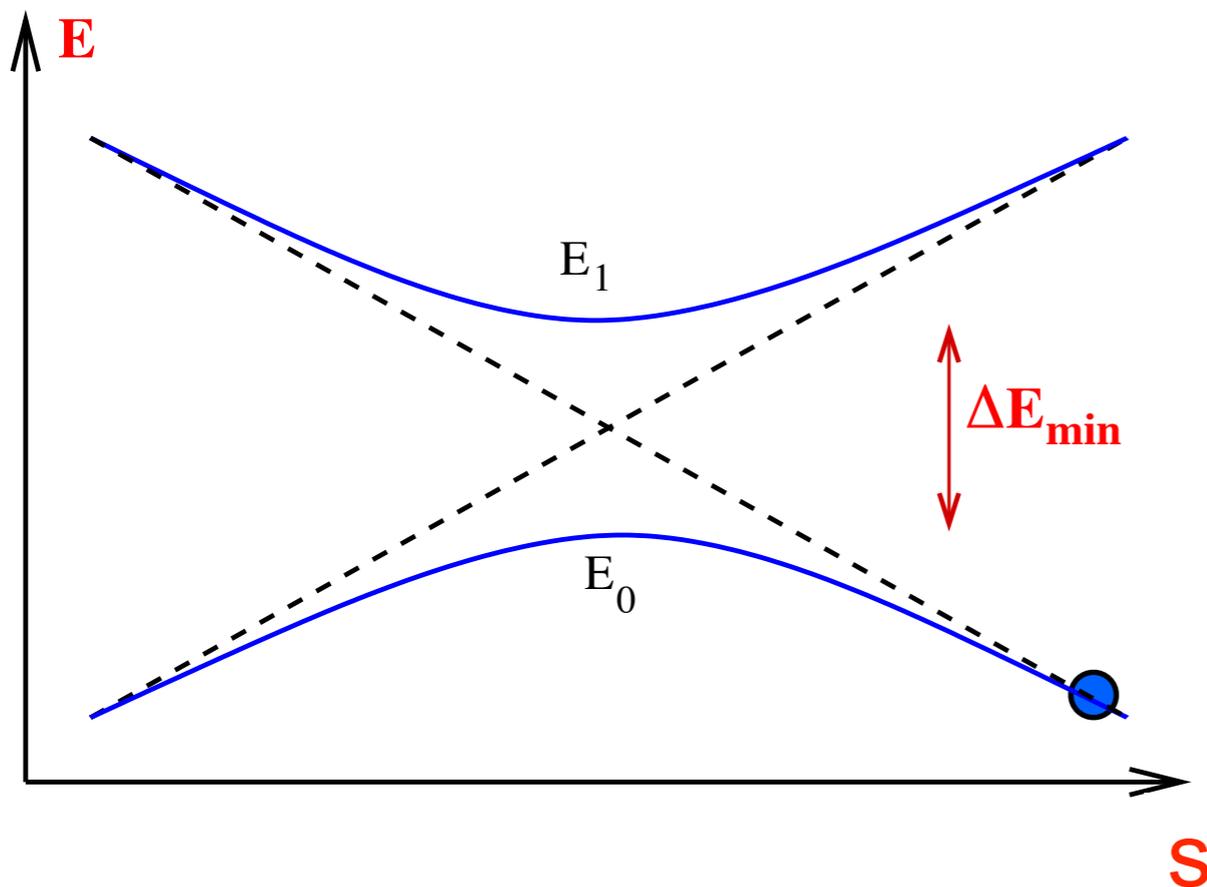


Landau Zener Theory:
To stay in the ground state the time needed is proportional to ΔE_{\min}^{-2}

Quantum Phase Transition



Bottleneck is likely to be a **quantum phase transition (QPT)** where the **gap to the first excited state is very small**

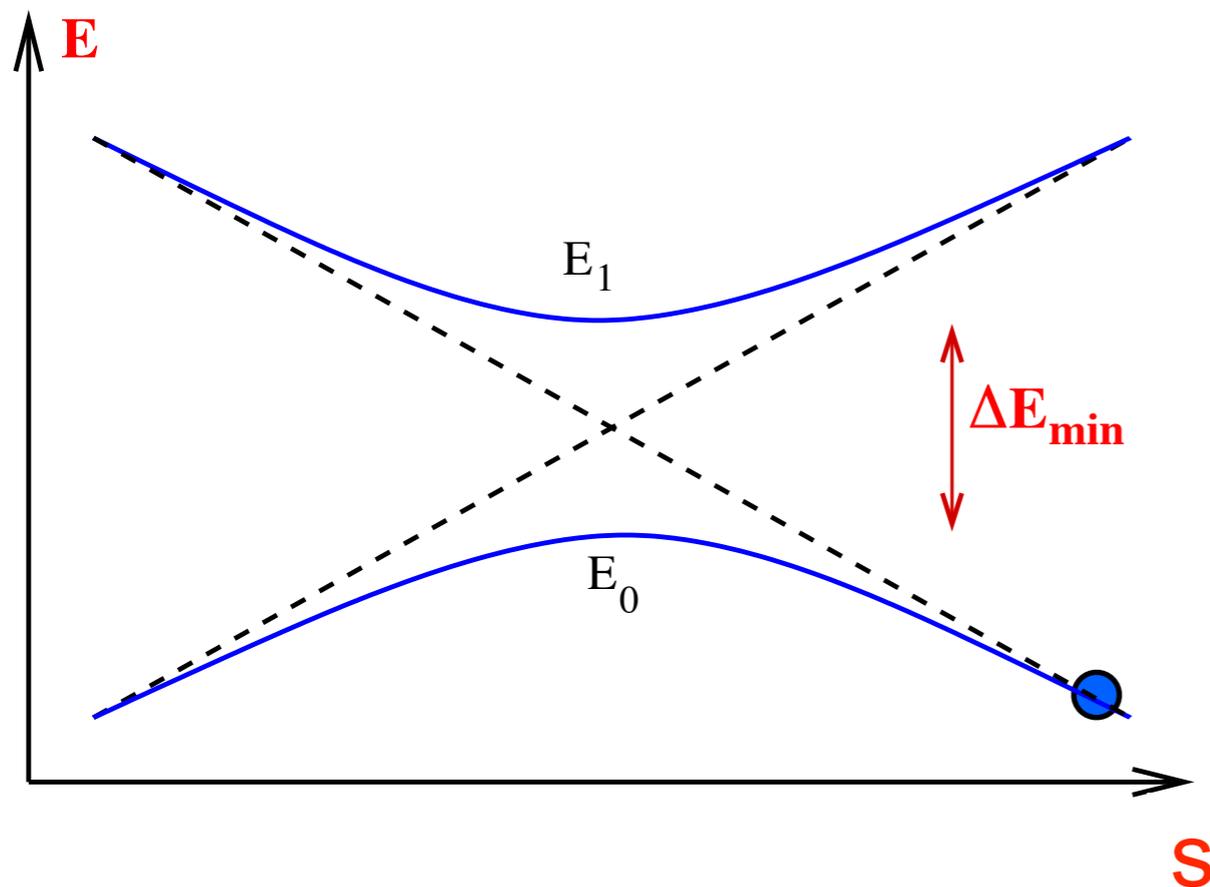


Landau Zener Theory:
To stay in the ground state the time needed is proportional to ΔE_{\min}^{-2}

Quantum Phase Transition



Bottleneck is likely to be a **quantum phase transition (QPT)** where the **gap to the first excited state is very small**



Landau Zener Theory:
To stay in the ground state the time needed is proportional to ΔE_{\min}^{-2}

Used **QMC** to compute ΔE for different **s**: $\rightarrow \Delta E_{\min}$

Synopsis: Mind the Gap



[iStockphoto.com/cinoby](https://www.iStockphoto.com/cinoby)

Exponential complexity of the quantum adiabatic algorithm for certain satisfiability problems

Itay Hen and A. P. Young

Phys. Rev. E **84**, 061152 (2011)

Published December 29, 2011

Quantum computers promise to accelerate some kinds of calculations in a remarkable manner. But as in present-day classical computing, hardware is only half the story: efficiency requires development of appropriate algorithms, such as the fast Fourier transform.

To apply a quantum computer to a broad class of problems, general-purpose algorithms are needed. One such method is the quantum adiabatic algorithm, in which the problem to be solved is coded into a Hamiltonian H . One prepares the quantum computer in the ground state of a reference Hamiltonian H_R and then has it evolve under a time-dependent Hamiltonian $H(t)$ that gradually switches from H_R to H . If the evolution is slow enough (“adiabatic”) the system ends up in the ground state of H , which contains information about the desired solution.

In a paper in *Physical Review E*, Itay Hen and Peter Young of the University of California, Santa Cruz, show that “slow enough” may be very slow indeed. The reason is that the time required for adiabatic evolution depends inversely on the gap in energies between the ground and first excited states of $H(t)$. Using computer simulations, Hen and Young show that for three classes of logic problems, the scaling of the gap is such that the computational time can be expected to grow exponentially with the size of the problem. The authors suggest that it might be possible to optimize the evolution of $H(t)$ to avoid the bottleneck associated with a vanishing gap. – *Ron Dickman*

Synopsis: Mind the Gap



iStockphoto.com/cinoby

Exponential complexity of the quantum adiabatic algorithm for certain satisfiability problems

Itay Hen and A. P. Young

Phys. Rev. E **84**, 061152 (2011)

Published December 29, 2011

Quantum computers promise to accelerate some kinds of calculations in a remarkable manner. But as in present-day classical computing, hardware is only half the story: efficiency requires development of appropriate algorithms, such as the fast Fourier transform.

To apply a quantum computer to a broad class of problems, general-purpose algorithms are needed. One such method is the quantum adiabatic algorithm, in which the problem to be solved is encoded into a Hamiltonian H . One prepares the quantum computer in the ground state of a reference time-dependent Hamiltonian $H(t)$ that gradually switches from H_R to H . If the evolution is slow enough, the system remains in the ground state of H , which contains information about the desired solution.

In a paper published in *Physical Review E*, Peter Young of the University of California, Santa Cruz, show that "slow enough" may be very slow indeed. The rate of evolution depends inversely on the gap in energies between the ground and first excited states of $H(t)$. Using a technique called the spectral gap method, they show that for three classes of logic problems, the scaling of the gap is such that the computational time can be exponentially large. The authors suggest that it might be possible to optimize the evolution of $H(t)$ to avoid the computational bottleneck.

MIND THE GAP

by Dickman

Synopsis: Mind the Gap



iStockphoto.com/cinoby

Quantum computers promise to accelerate some kinds of calculations, but that's only half the story: efficiency requires development of new algorithms.

To apply a quantum computer to a broad class of problems, researchers have developed a new algorithm, in which the problem to be solved is encoded into the ground state of a Hamiltonian.

Hamil
("adia

In a p
The r
Using
be ex
bottle

MIND THE GAP

Exponential complexity of the quantum adiabatic algorithm for certain satisfiability problems

Itay Hen and A. P. Young

PH
Pu



Peter
ev
ow th
of th
on D

d.
an
ne

Quantum Monte Carlo

Early numerics, Farhi et al. for very small sizes $N \leq 20$, on a particular problem found the time varied only as N^2 , i.e.

polynomial!

But possible “crossover” to exponential at larger sizes?

Want to estimate how the running time varies with size **for large sizes.**

We therefore have to do a **Quantum Monte Carlo Sampling** of the 2^N states.

QMC can efficiently study only **equilibrium** properties of a quantum system by simulating a **classical model with an extra dimension**, imaginary time, τ , where $0 \leq \tau < 1/T$.
Used a version called the stochastic series expansion (SSE), pioneered by Sandvik.

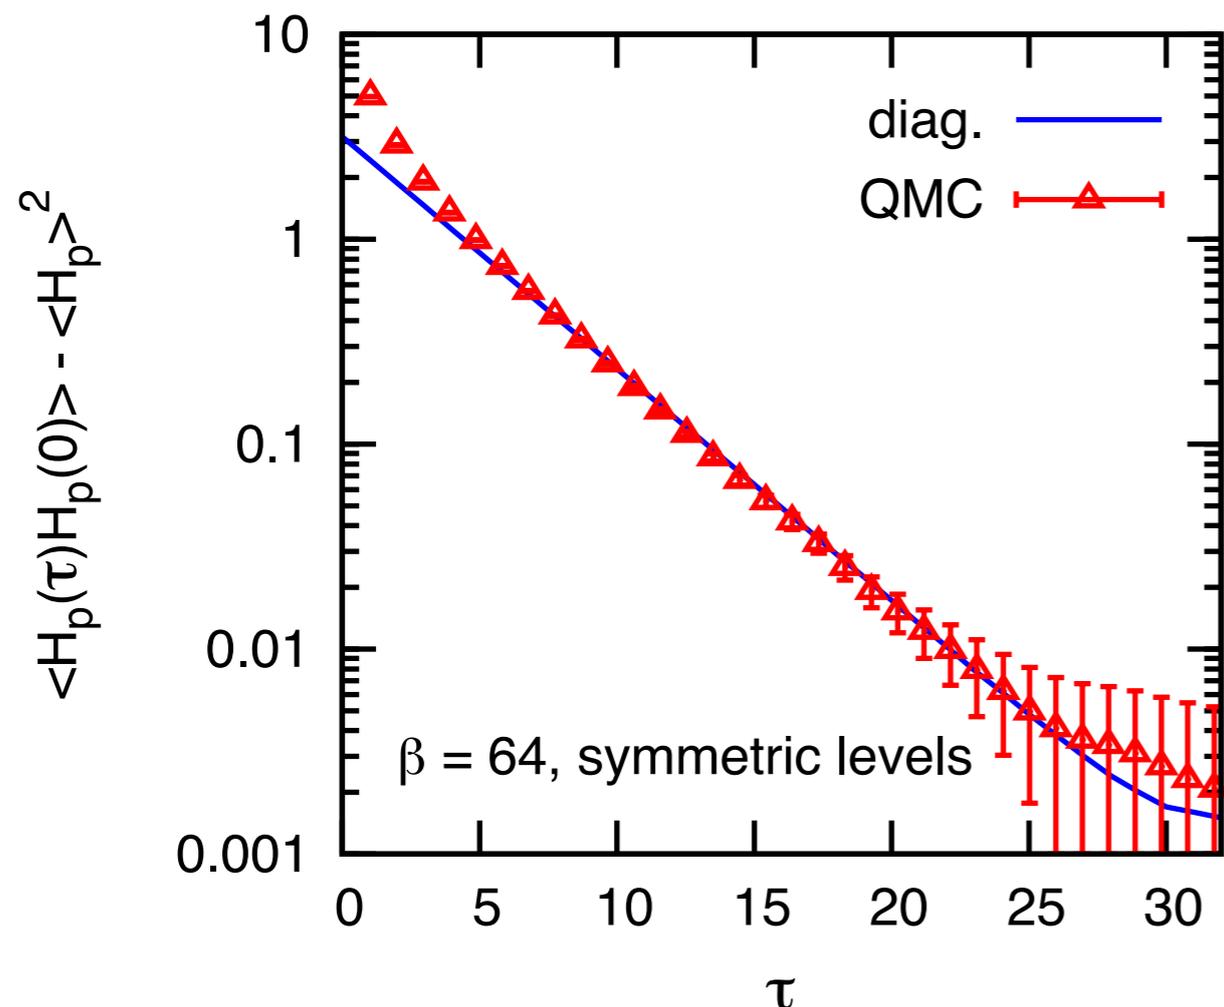
Not perfect, (statistical errors, need to ensure equilibration) but the only numerical method available for large N.

Examples of results with the SSE code

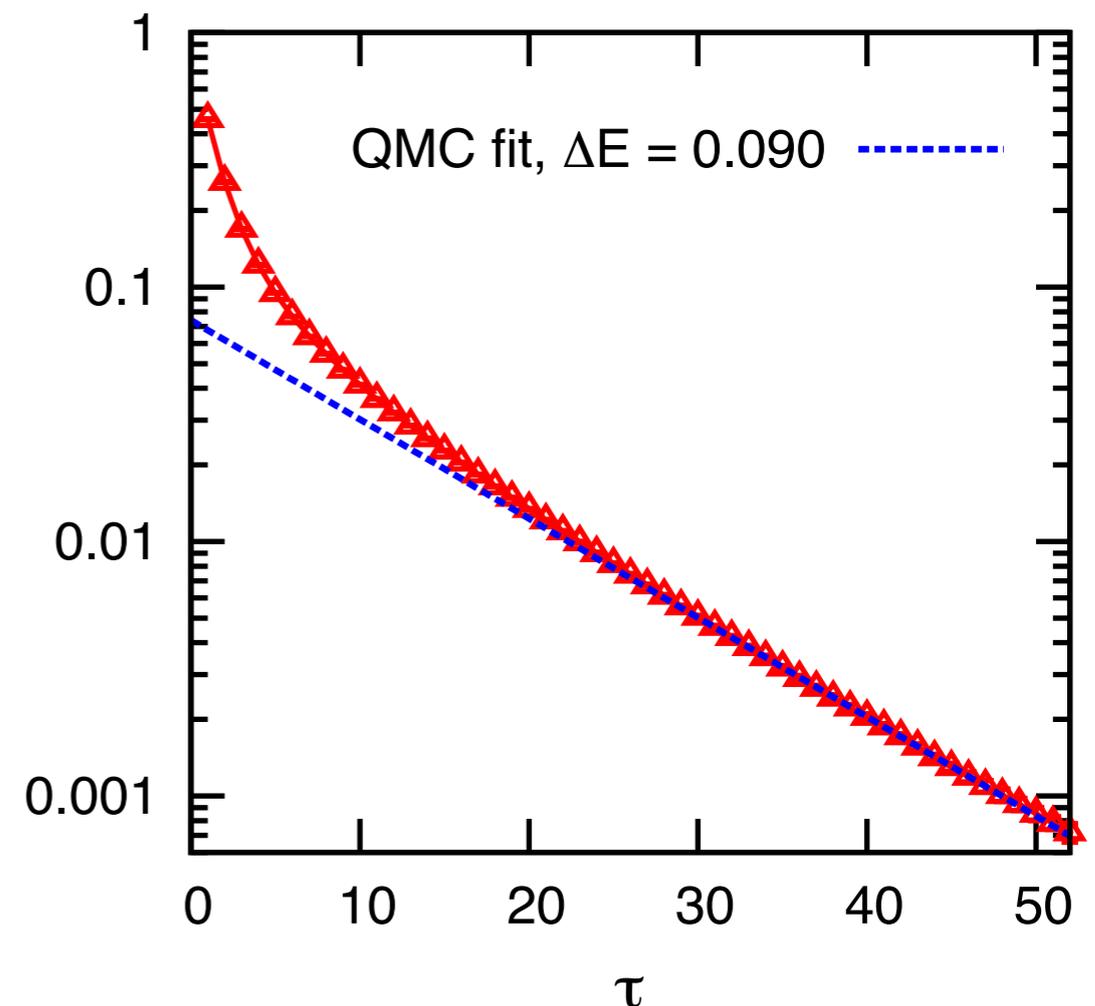
Time dependent correlation functions decay with τ as a **sum of exponentials**

$$\langle A(\tau)A(0) \rangle - \langle A \rangle^2 = \sum_{n \neq 0} |\langle 0|A|n \rangle|^2 \exp[-(E_n - E_0)\tau]$$

For **large τ** only first excited state contributes, \rightarrow **pure exponential decay**

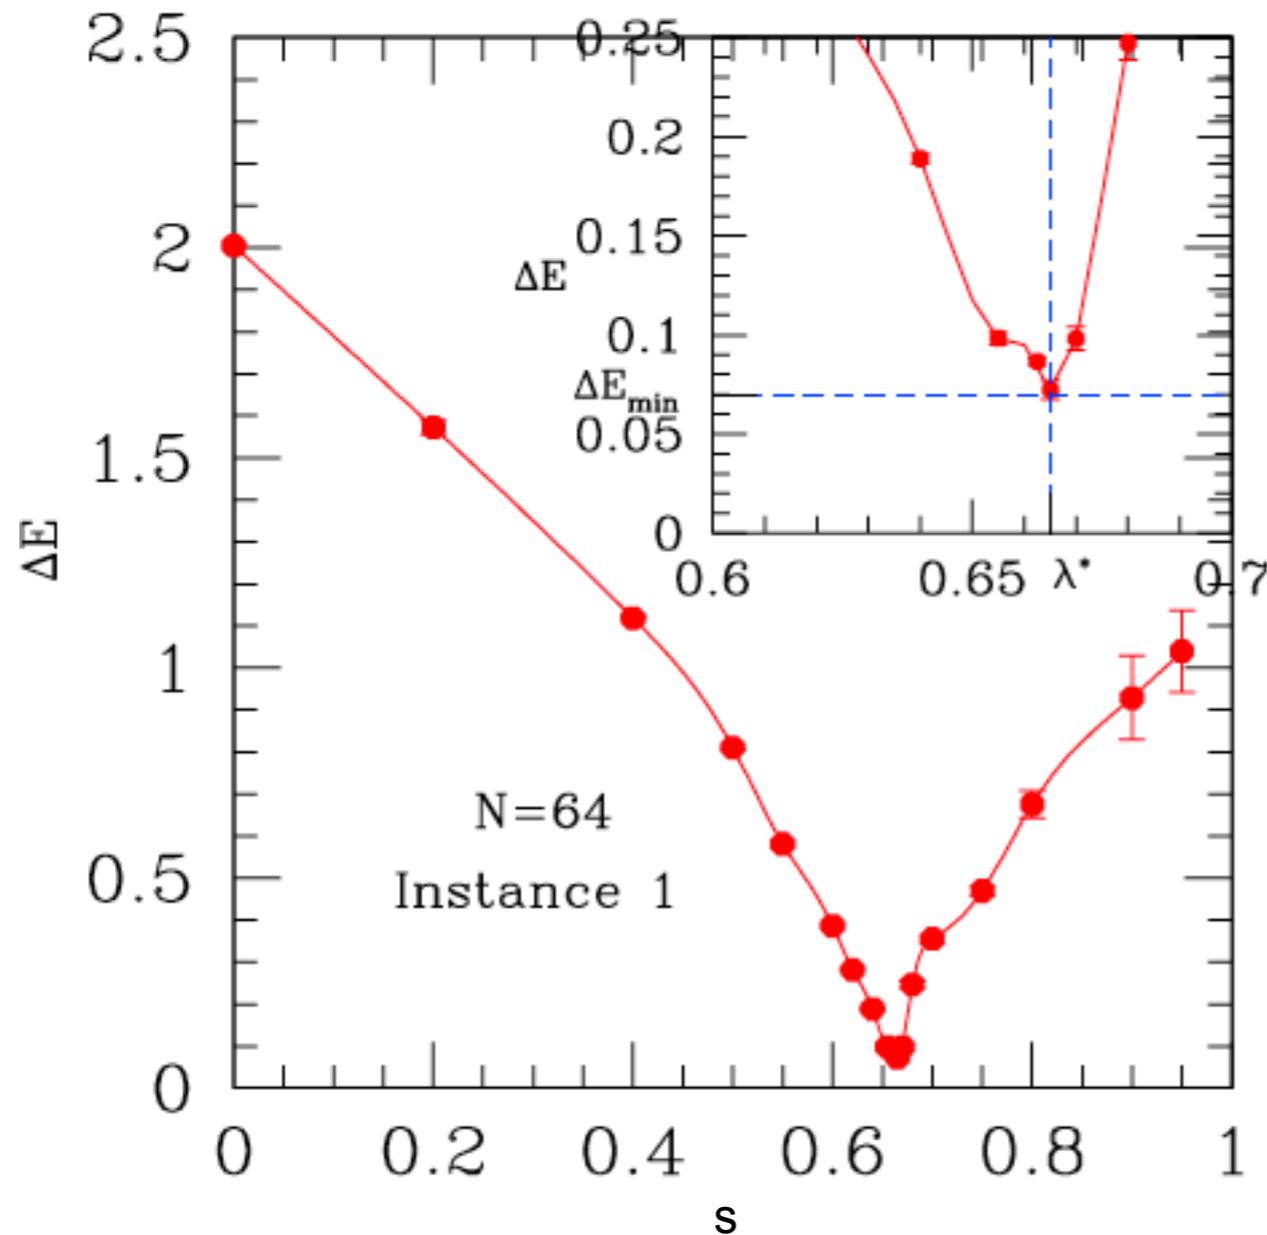


Small size, N = 24, excellent agreement with diagonalization.



Large size, N = 128, good quality data, slope of straight line \rightarrow gap.

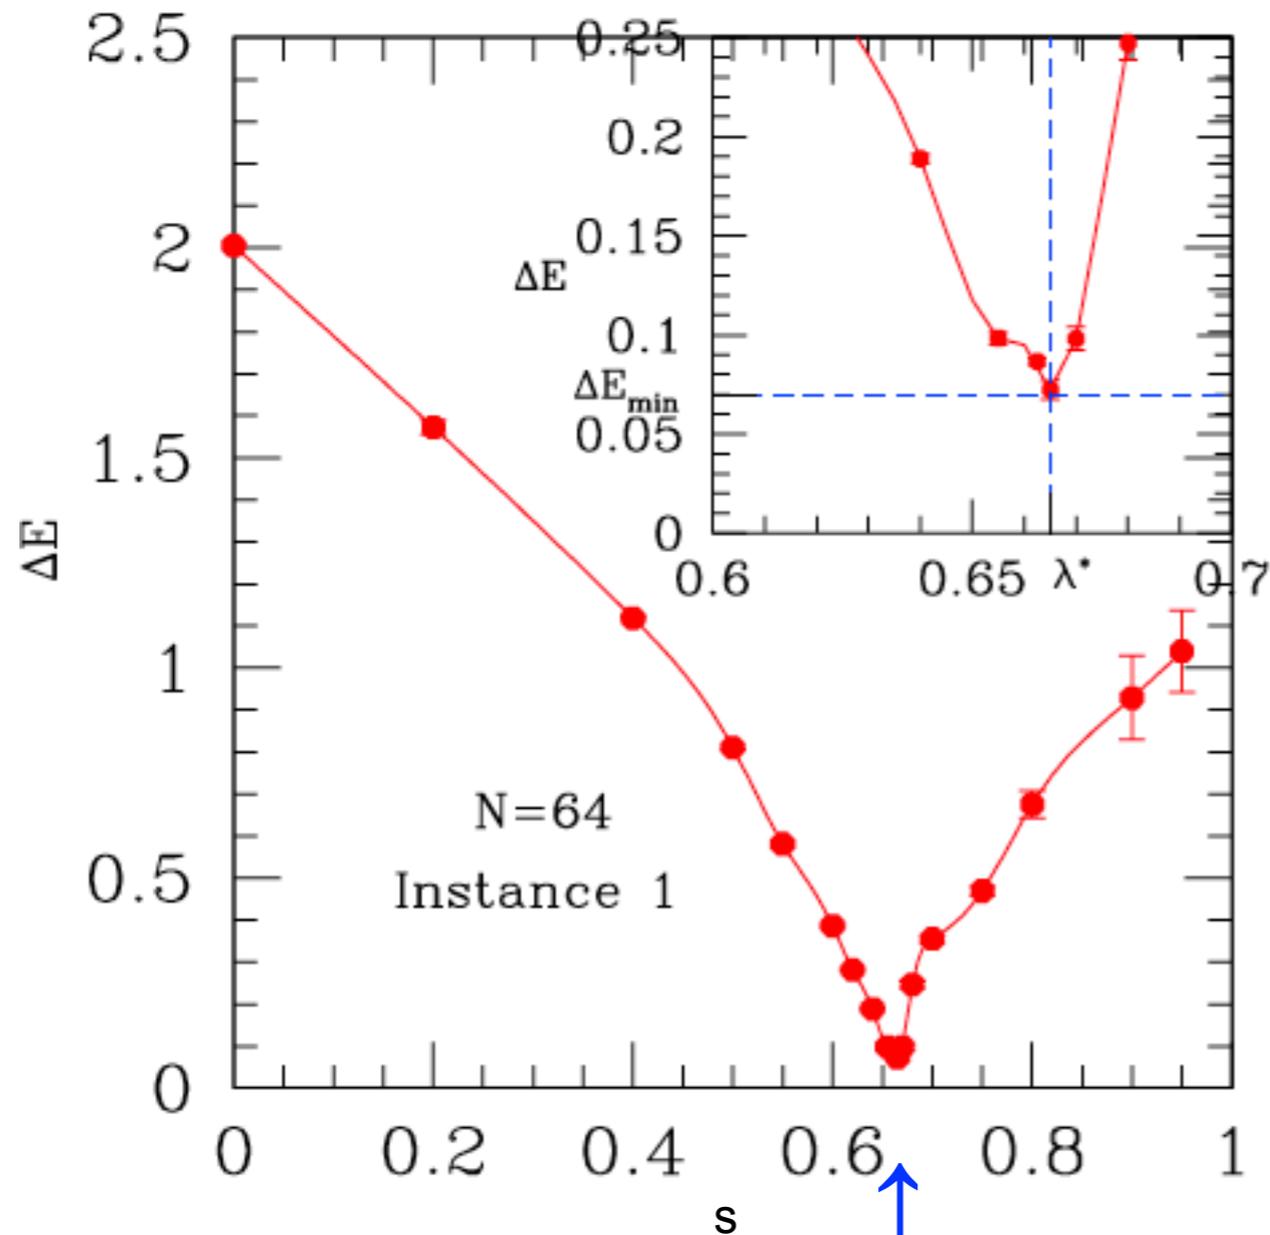
Dependence of gap on s



Results for the dependence of the gap to the first excited state, ΔE , with s , for one instance of 1-in-3 SAT with $N = 64$.

The gap has a minimum for s about 0.66 which is the bottleneck for the QAA.

Dependence of gap on s

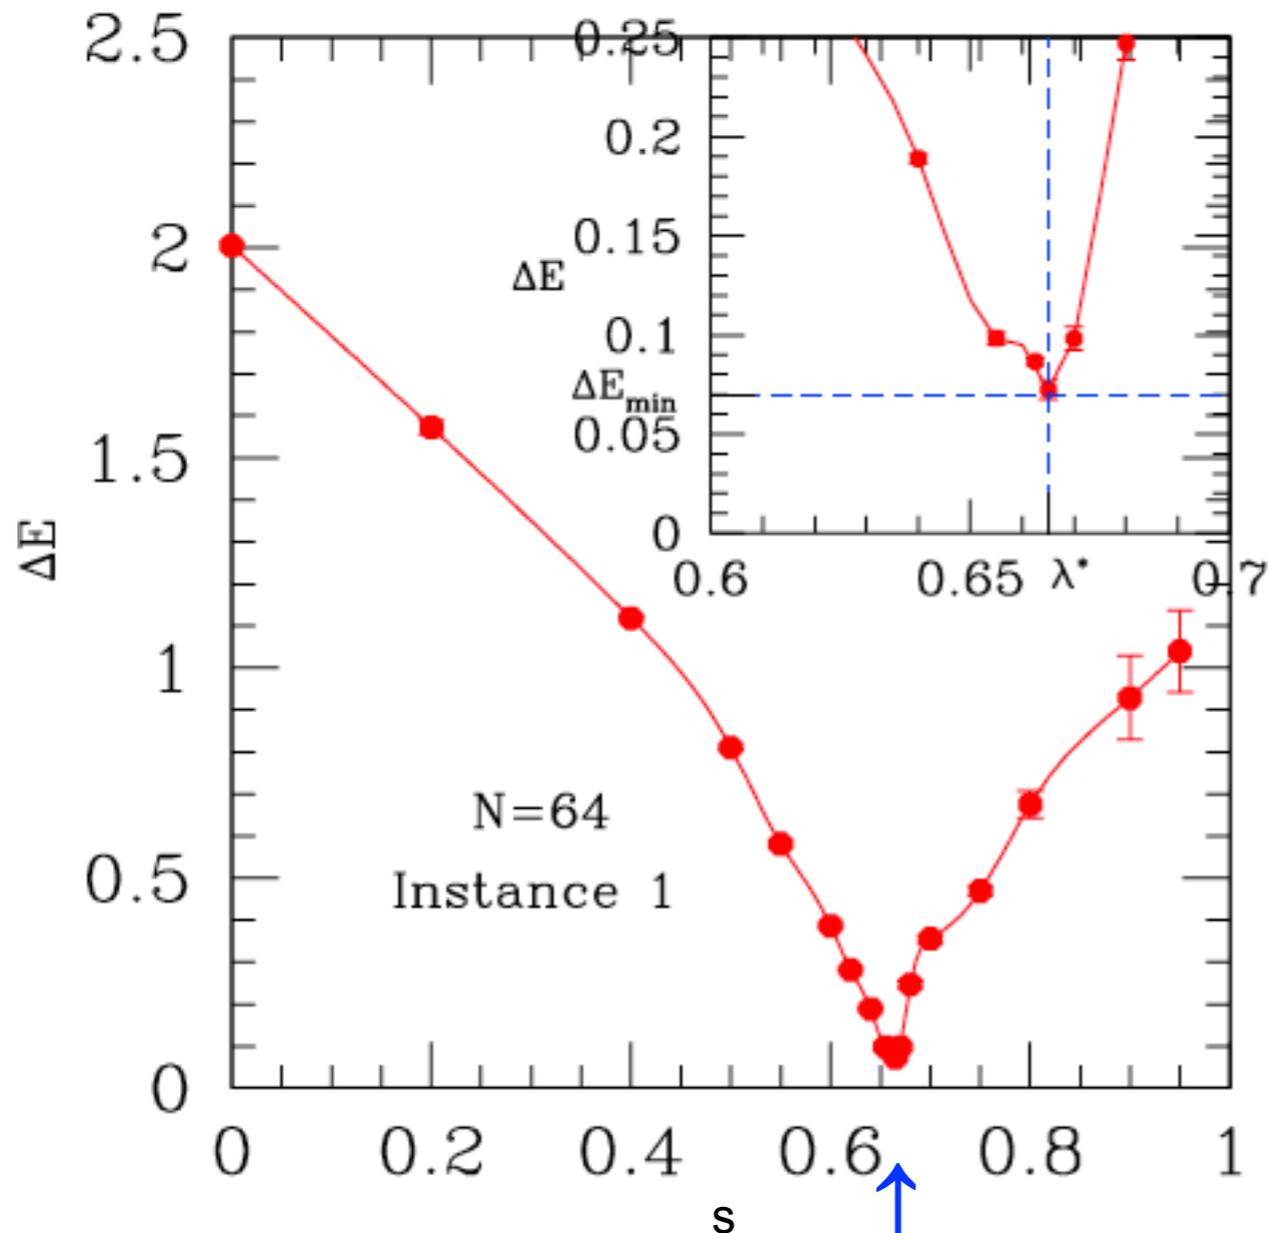


Mind THIS gap

Results for the dependence of the gap to the first excited state, ΔE , with s , for one instance of 1-in-3 SAT with $N = 64$.

The gap has a minimum for s about 0.66 which is the bottleneck for the QAA.

Dependence of gap on s



Mind THIS gap

Results for the dependence of the gap to the first excited state, ΔE , with s , for one instance of 1-in-3 SAT with $N = 64$.

The gap has a minimum for s about 0.66 which is the bottleneck for the QAA.

We compute the minimum gap for many (50) instances for each size N and look how the median minimum gap varies with size.

Results

We studied several models using Quantum Monte Carlo

- A **spin glass**. (Not the usual spin glass model though.)
- Several **satisfiability (SAT)** problems

We do not include the effects of finite-T and non-thermal noise in the D-Wave machine.

Our aim is to see if there is a quantum speedup in the QA/QAA paradigm of quantum computing in an ideal world where these complications do not occur.

A spin glass

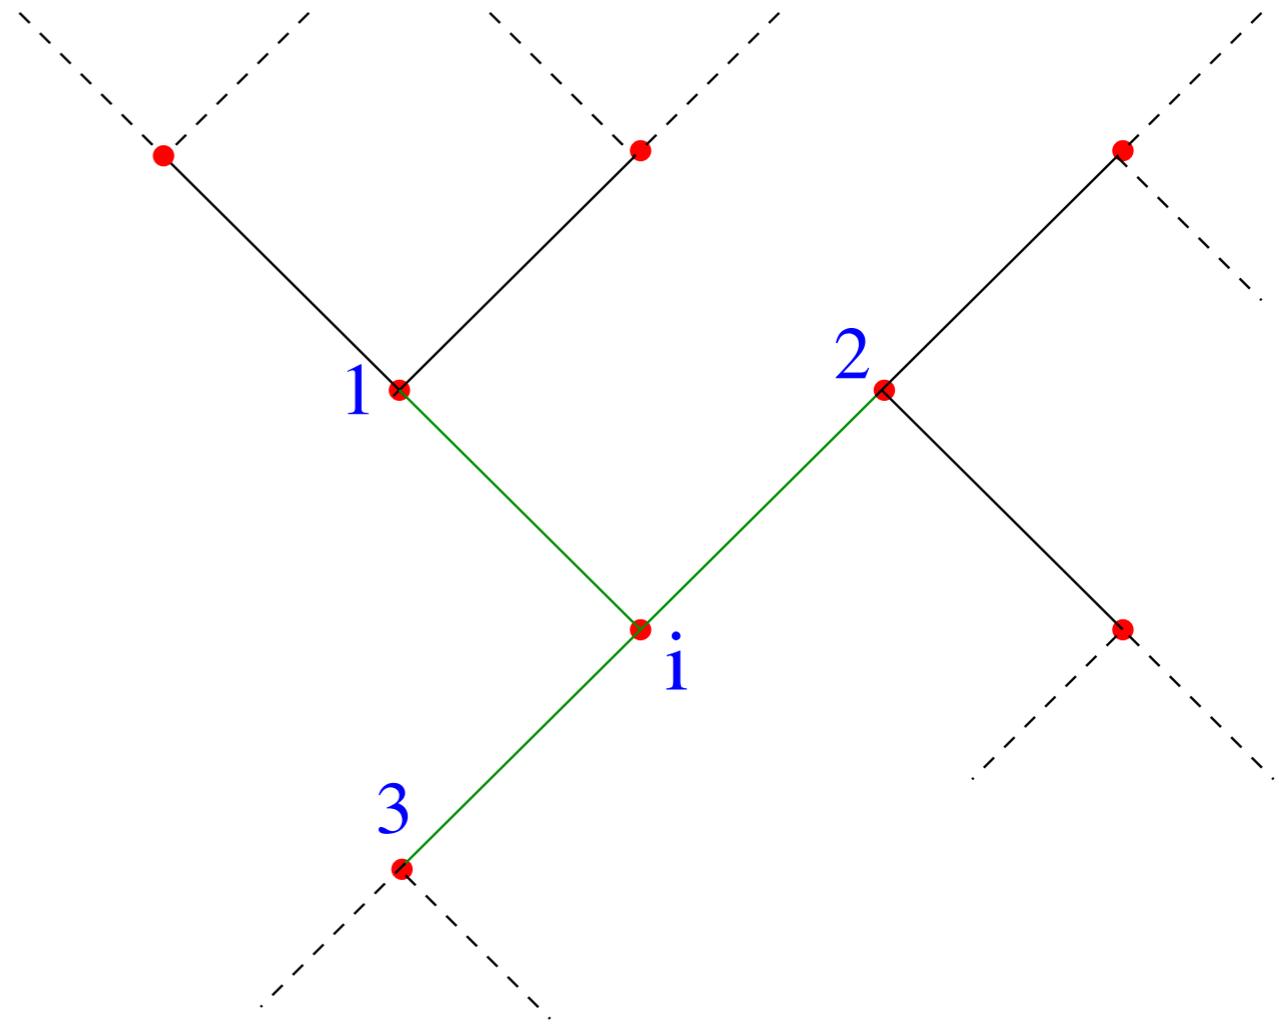
We study a very simple spin glass model, an **anti-ferromagnet** on a **regular random graph**.

$$\mathcal{H}_P = \frac{1}{2} \sum_{\langle i,j \rangle} \left(1 + \sigma_i^z \sigma_j^z \right)$$

3-Regular means that each spin is connected to exactly three neighbors.

Spin glasses need **randomness and frustration**. Here we have no disorder in the interactions so these come entirely from the geometry of the (large) loops.

Because of large loops there is **no AF state, only a spin glass** (Zdeborova et al).



Note: there are large loops

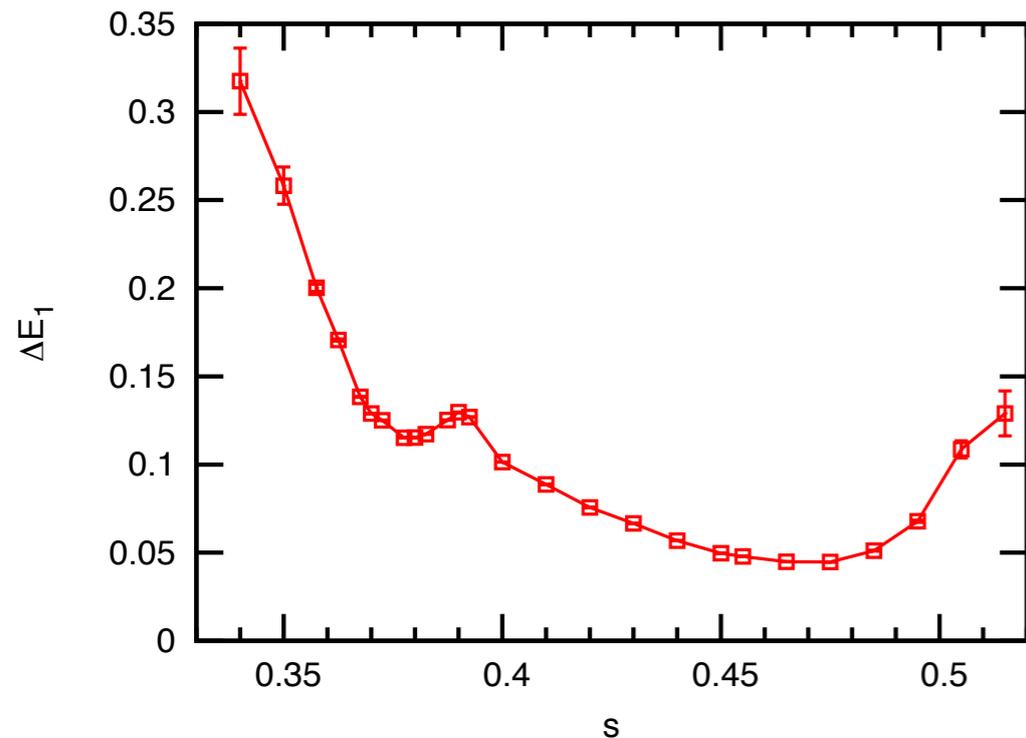
Spin glass

Adding the driver Hamiltonian there is a quantum phase transition at $s = s^*$ above which the symmetry is spontaneously broken. This is the quantum spin glass phase.

“Cavity” calculations (Gosset, Zamponi) find $s^* \cong 0.36$
Investigated the problem near s^* ($s \leq 0.5$).

Also just considered instances with a “unique satisfying assignment” (apart from the degenerate state related by flipping all the spins). (These are exponentially rare.)

Spin glass on a random graph: (Gap)



For larger sizes, a fraction of instances have **two minima**, one fairly close to s^* (≈ 0.36) and other at larger s in the spin glass phase.

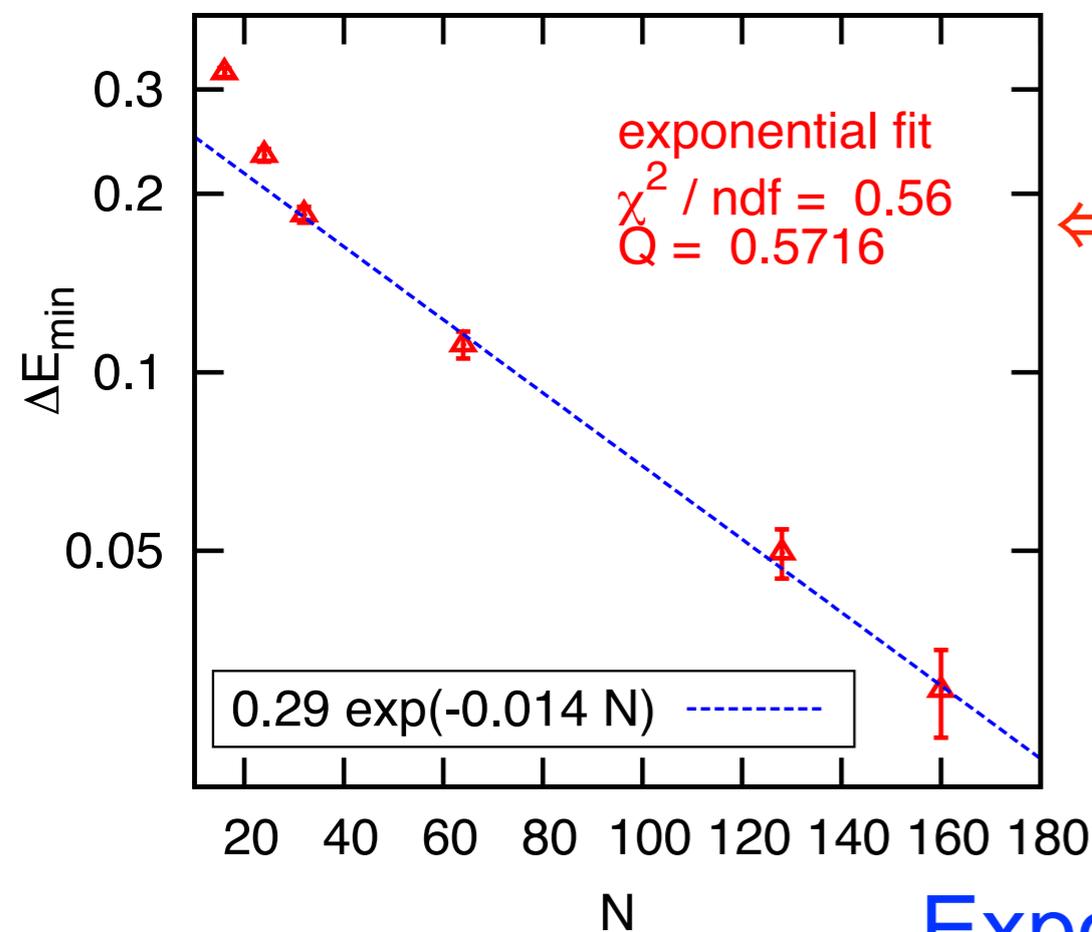
← **Figure** shows an example for $N = 128$.

Hence did **2 analyses**

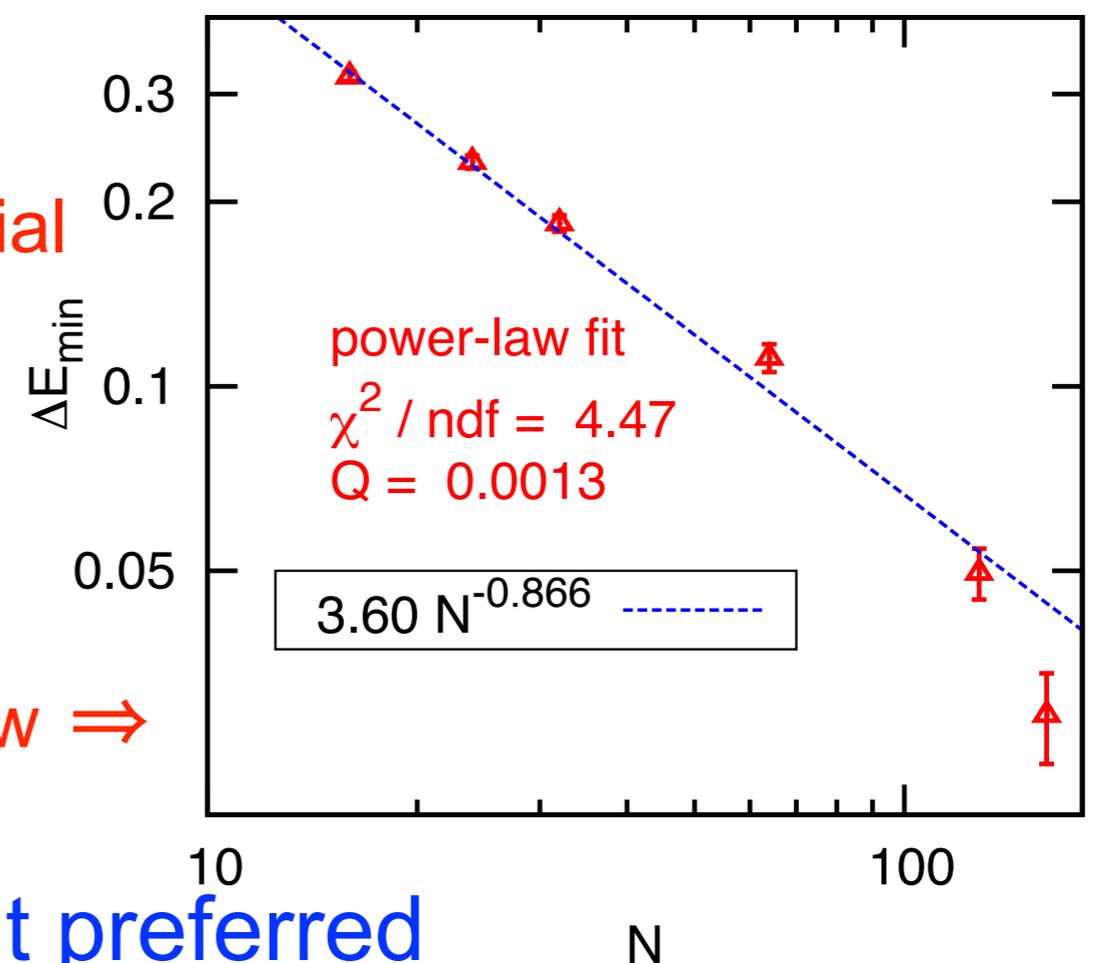
(i) **Global** minimum in range (up to $s=0.5$)

(ii) If two minima, just take the **local** minimum near s^* .

Figures below show the fits to the **global data**



← **exponential**



power law ⇒

Exponential fit preferred

Summary for the spin glass

In the spin glass phase the minimum gap decreases exponentially

We also found that, at the quantum critical point, the minimum gap seems to decrease polynomially.

Hence the bottleneck for the spin glass is in the quantum spin glass phase.

This is different from the SAT problems where the bottleneck is AT the phase transition.

SAT problems

We also studied some Satisfiability (**SAT**) problems, very popular with computer scientists.

One has **N** bits, and **M** logical conditions (**clauses**) each of which involves a small number of bits.

Question: is there a configuration of the bits which satisfies all the clauses?

If **M** \ll **N** then **yes** (many ways to do it)

If **M** \gg **N** then **no** (always a conflict)

Hence there is a **satisfiability transition** at a critical value of **M/N**. The problem is hard in this region, so we shall **take M/N close to critical value**.

Convert to a problem Hamiltonian **H_P** which is a sum of contributions from each clause:

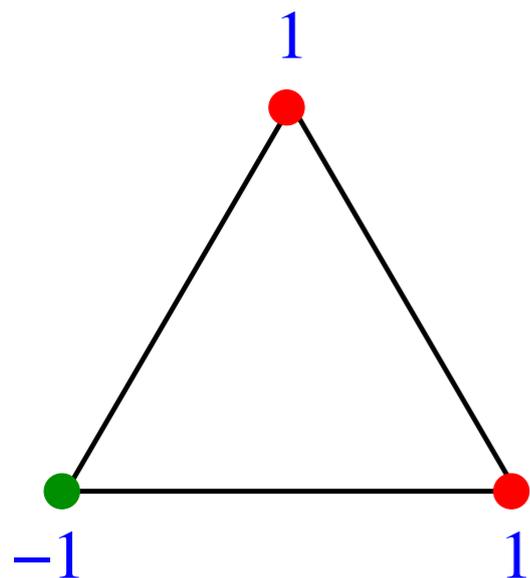
$$\mathcal{H}_P = \sum_{a=1}^M \mathcal{H}_a$$

H_P is chosen such that it is zero if the clause is satisfied and a positive value if it is not. **Hence the problem is SAT if the ground state energy is zero.**

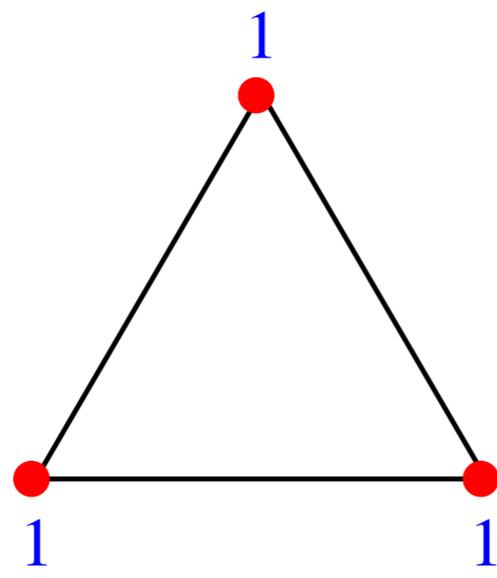
Example of a Satisfiability Problem

- 1-in-3 SAT

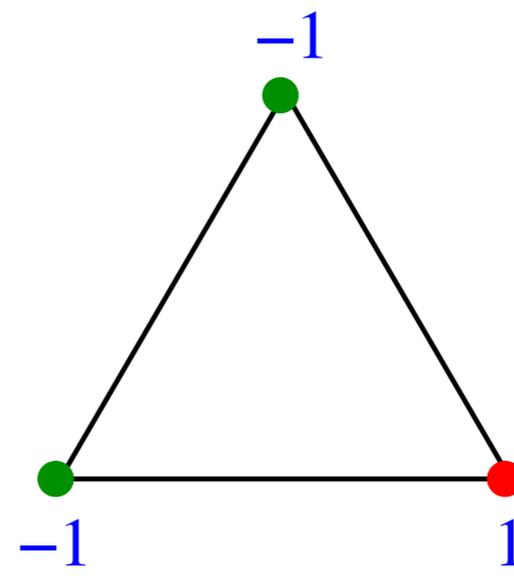
The clause is formed from 3 bits picked at random. The clause is satisfied (has energy 0) if one is 1 and the other two are 0 (in terms of spins one is -1 (green) and the other two are +1 (red)). Otherwise it is not satisfied (the energy is 1).



Satisfied



Unsatisfied

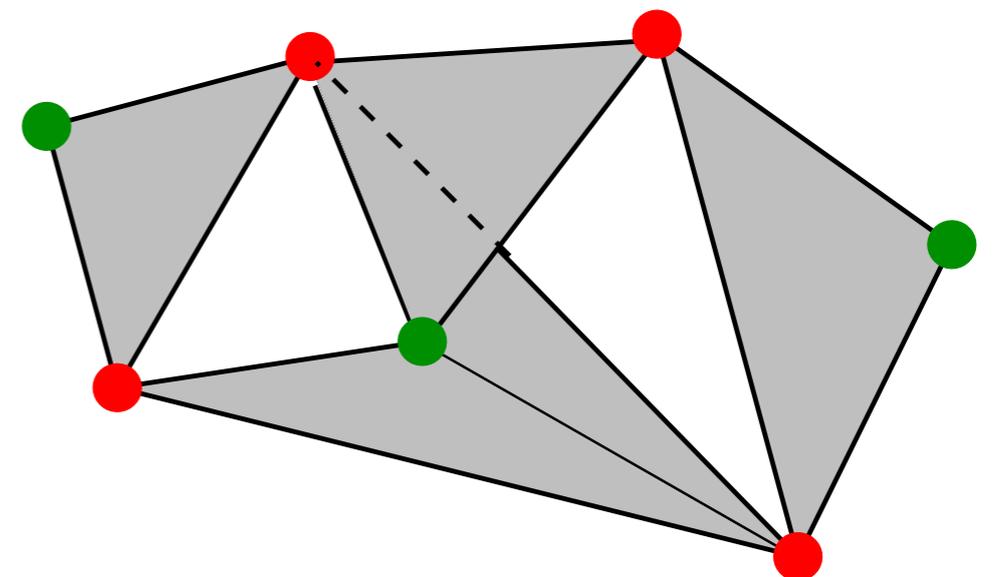


Unsatisfied

Example of a satisfying assignment with $N=7$, $M=5$.

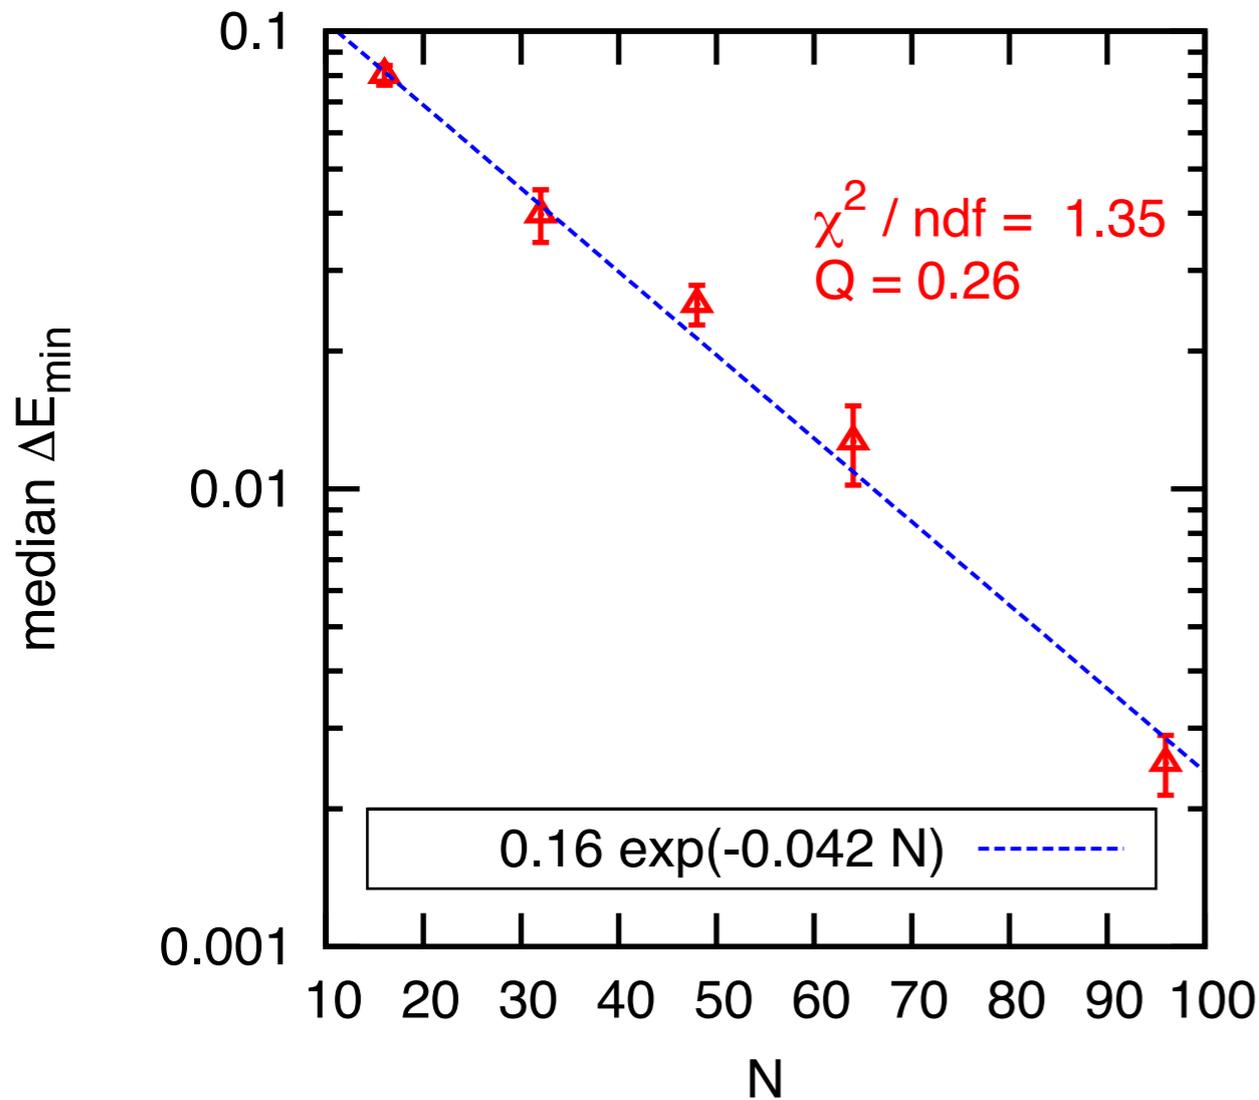
(V. Choi)

$$\mathcal{H}_P = \sum_{\text{clauses}} \left(\frac{\sigma_1^z + \sigma_2^z + \sigma_3^z - 1}{2} \right)^2$$

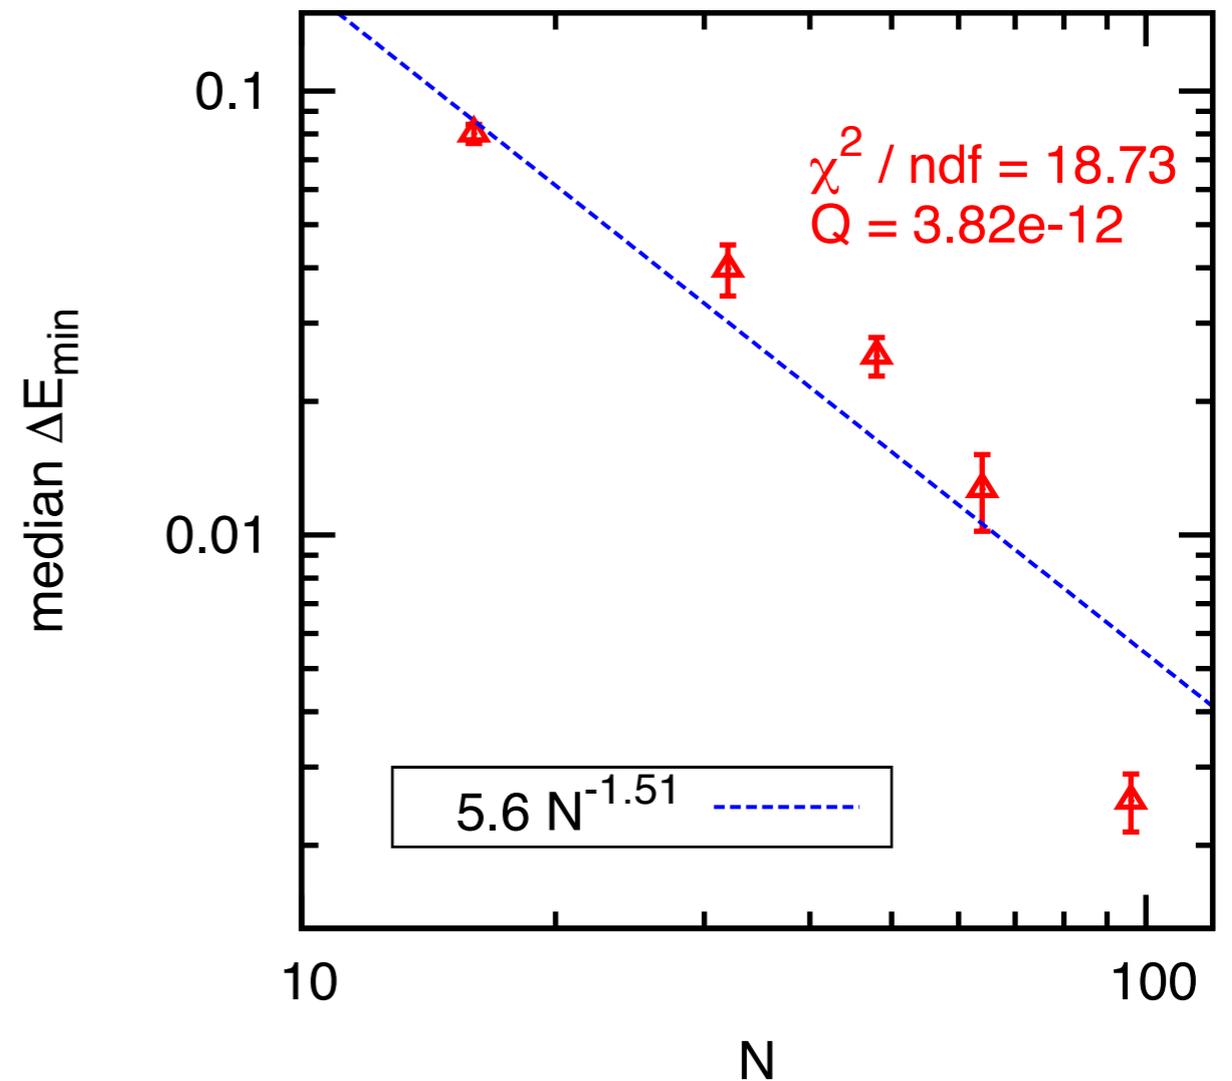


1-in-3 SAT (gap)

Plots of the median minimum gap (average over 50 instances)



Exponential fit

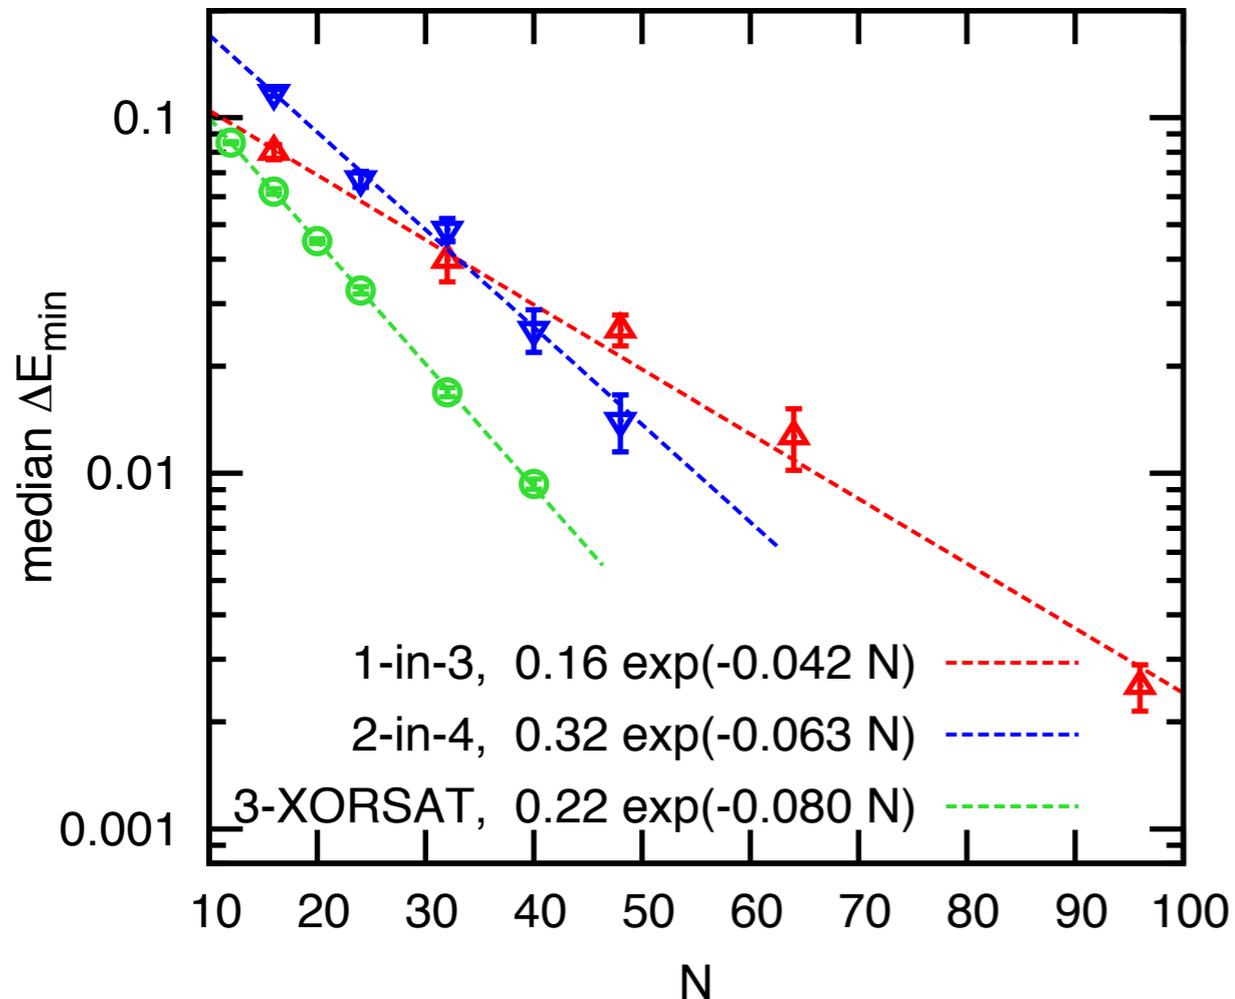


Power law fit

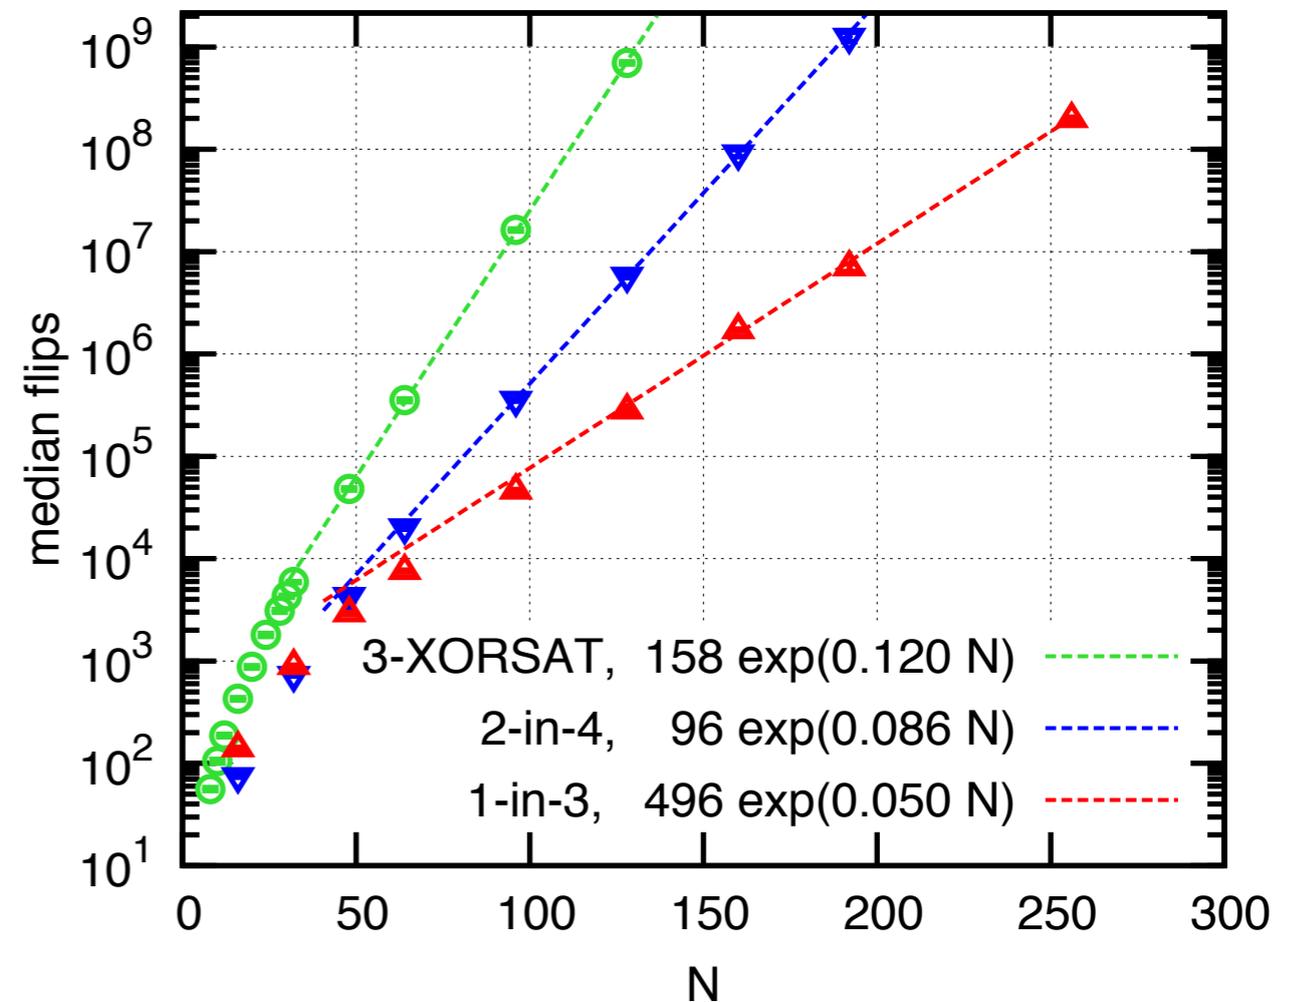
Clearly the behavior of the minimum gap is exponential

3-SAT models: comparison between QAA and a classical algorithm, WalkSAT

QAA



WalkSAT



WalkSAT is a classical heuristic algorithm similar in spirit to SA.
Exponential behavior for both QAA and WalkSAT

Problems that are hardest for WalkSAT are hardest for QAA and vice-versa

Curious that the **hardest** problem for these heuristic algorithms is the one with a **polynomial time** algorithm (complexity class **P**).

Summary for SAT models

For **WalkSAT** we find that the number of “flips” is given by

$$\tau \propto \exp(\mu N)$$

and we determine the values for μ .

For the **QAA** we find $\Delta E_{\min} \propto \exp(-c N)$

and we determine the values for c . Running time $\propto 1/(\Delta E_{\min})^2$

so we have $\mu = 2c$ and hence we can compare the values of μ for the classical and quantum algorithms:

model	$\mu(\text{QAA})$	$\mu(\text{WalkSAT})$	Ratio
1-in-3 SAT	0.084(3)	0.0505(5)	1.66
2-in-4 SAT	0.126(5)	0.0858(8)	1.47
3-XORSAT	0.159(2)	0.1198(20)	1.32

QAA is always a bit worse (i.e. μ a bit bigger) than WalkSAT.

Bottleneck for SAT is AT the (first order) quantum phase transition. Transition first order argued by Krzakala et al.

Conclusions

Conclusions

- We have studied the **efficiency of the QAA** neglecting the effects of non-zero temperature and non-thermal noise which are present in the D-Wave machine.

Conclusions

- We have studied the **efficiency of the QAA** neglecting the effects of non-zero temperature and non-thermal noise which are present in the D-Wave machine.
- Even in this ideal world we do **NOT find the QAA to be more efficient than an analogous, heuristic, classical algorithm, WalkSAT.**

Conclusions

- We have studied the **efficiency of the QAA** neglecting the effects of non-zero temperature and non-thermal noise which are present in the D-Wave machine.
- Even in this ideal world we do **NOT find the QAA to be more efficient than an analogous, heuristic, classical algorithm, WalkSAT.**
- We have taken the **simplest implementation** of the QAA with all the transverse fields and weights of the clauses (for SAT) equal. One can try to change these, either “on the fly” or by brute force.

Conclusions

- We have studied the **efficiency of the QAA** neglecting the effects of non-zero temperature and non-thermal noise which are present in the D-Wave machine.
- Even in this ideal world we do **NOT find the QAA to be more efficient than an analogous, heuristic, classical algorithm, WalkSAT.**
- We have taken the **simplest implementation** of the QAA with all the transverse fields and weights of the clauses (for SAT) equal. One can try to change these, either “on the fly” or by brute force.
- Also interest in using QAA to solve other problems such as **graph isomorphism.**

Conclusions

- We have studied the **efficiency of the QAA** neglecting the effects of non-zero temperature and non-thermal noise which are present in the D-Wave machine.
- Even in this ideal world we do **NOT find the QAA to be more efficient than an analogous, heuristic, classical algorithm, WalkSAT.**
- We have taken the **simplest implementation** of the QAA with all the transverse fields and weights of the clauses (for SAT) equal. One can try to change these, either “on the fly” or by brute force.
- Also interest in using QAA to solve other problems such as **graph isomorphism.**
- Other work (Troyer et al.) has compared results from the D-Wave machine with results from simulated annealing (SA) and simulated quantum annealing (SQA). They find the D-Wave machine **is** doing quantum annealing, but does **not appear to be more efficient than a classical algorithm.**

Conclusions

- We have studied the **efficiency of the QAA** neglecting the effects of non-zero temperature and non-thermal noise which are present in the D-Wave machine.
- Even in this ideal world we do **NOT find the QAA to be more efficient than an analogous, heuristic, classical algorithm, WalkSAT.**
- We have taken the **simplest implementation** of the QAA with all the transverse fields and weights of the clauses (for SAT) equal. One can try to change these, either “on the fly” or by brute force.
- Also interest in using QAA to solve other problems such as **graph isomorphism.**
- Other work (Troyer et al.) has compared results from the D-Wave machine with results from simulated annealing (SA) and simulated quantum annealing (SQA). They find the D-Wave machine **is** doing quantum annealing, but does **not appear to be more efficient than a classical algorithm.**

The latter is in agreement with us.