Quantum Annealing: Challenges & Prospects

Tameem Albash
Information Sciences Institute @USC
29 November 2018

From Bits to Qubits

Goal is to solve an inherently classical problem
with a quantum algorithm

Classical variables will be binary (bits)

$$x = \{0, 1\}$$

Quantum algorithm promotes bits to qubits

$$|\psi\rangle = a|0\rangle + b|1\rangle$$

A natural binary valued system is the Ising spin

$$\uparrow \equiv 0 \qquad \downarrow \equiv 1$$

Pick a basis for the qubits to associate classical bit values to

Computational basis:
eigenstates of the
Pauli-Z operator σ^z

$$\sigma^z\,|0\rangle = +\,|0\rangle \qquad \sigma^z\,|1\rangle = -\,|1\rangle$$

Our Objective Today

Find the ground state of Ising Hamiltonians

$$H_{\text{Ising}} = \sum_i h_i \sigma_i^z + \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z$$

Or any Hamiltonian defined entirely in terms of σ^z
(diagonal in the computational basis)

Local fields h_i and Ising couplings J_{ij} specify the problem

Spins live on the vertices of the connectivity graph

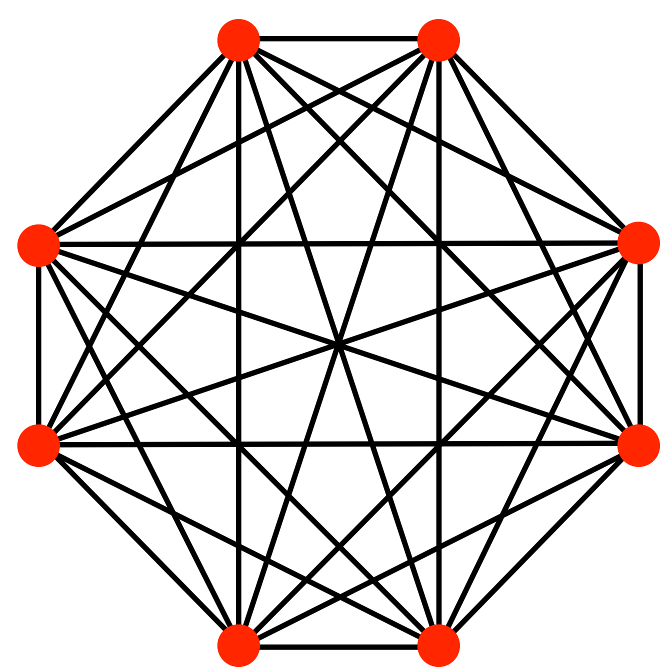

Weighted edges of the graph correspond to Ising interaction

Total of 2^N configurations for an N spin problem
Computationally prohibitive to search for the ground state

Ising is Sufficient

Algorithms to solve the Ising model is an important line of research

Physics

Classic system for studying magnetization, phase transitions, glassiness

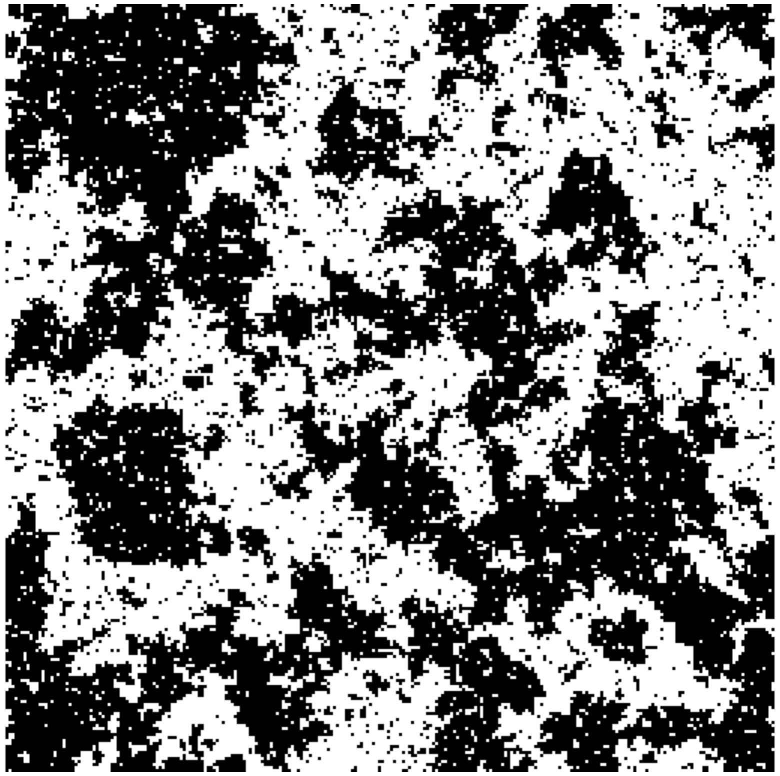

Computer science

The Ising problem on non-planar graphs belongs to the complexity class NP-hard[1]

A very good algorithm for solving the Ising model is likely to be a good algorithm for solving problems in NP

(1) F. Barahona, J. Phys. A: Math. Gen. 15 3241 (1982)

Why Solving NP Problems is Important

NP problems are ubiquitous with with range of applications

NP 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, software V&V
0-1 Integer Linear Programming	Natural language processing
Sub-graph isomorphism	Chem-informatics, drug discovery
Job shop scheduling	Manufacturing
MAX-2SAT	Artificial intelligence

Optimum solution to optimization problem

Finding ground state to Ising problem

These problems are typically very hard to solve,
requiring a time that grows exponentially with the problem size

Can Quantum Computing Help?

Can quantum computing help solve
a class of Ising problems more efficiently?

Quantum annealing (QA)[1,2]
is an adiabatic-paradigm quantum algorithm
to solve for the ground state of Ising problems

Good news

Provable speedups for oracular problems[3-5]

→

Bad news

Hamiltonians involve N-body operators

No proof that a speedup is impossible

←

No provable or demonstrated speedups for Ising-like Hamiltonians with bounded locality

(1) T. Kadowaki, and H. Nishimori , Phys. Rev. E 58 (5), 5355 (1998)
(2) E. Farhi,et al., Science 292 (5516), 472 (2001)
(3) J. Roland and N. Cerf, Phys. Rev. A 65, 042308 (2002)
(4) I. Hen, Europhysics Letters 105 (5), 50005 (2014)
(5) R. D. Somma, et al. , Phys. Rev. Lett. 109 (5), 050501 (2012)

Outline

1. How should QA work?

Closed-system QA

2. Why has QA not worked?

Open-system QA and the D-Wave processors

3. How do we address these challenges?

Beyond standard QA

What is Standard Quantum Annealing?

A continuous interpolation
between two Hamiltonians

H_0
Easily prepared
ground state

$$H_0 = -\sum_i \sigma_i^x$$

$$H(t) = \left(1 - \frac{t}{t_f}\right) H_0 + \frac{t}{t_f} H_1$$

$[H_0, H_1] \neq 0$

$t \in [0, t_f]$

H_1
Ground state
encodes solution

$$H_1 = H_{\text{Ising}}$$

Procedure for Quantum Annealing (QA):

(1)
Prepare system in the ground state of $H(t=0)$

(2)
Evolve according to $H(t)$ for a time t_f

(3)
Measure the state in the σ^z basis

Success if measurement outcome is the ground state of H_1

A Guarantee for QA to Find the Solution

Adiabatic theorem

if the interpolation is sufficiently slow
then with high probability the final state of the system is the ground state of H_1

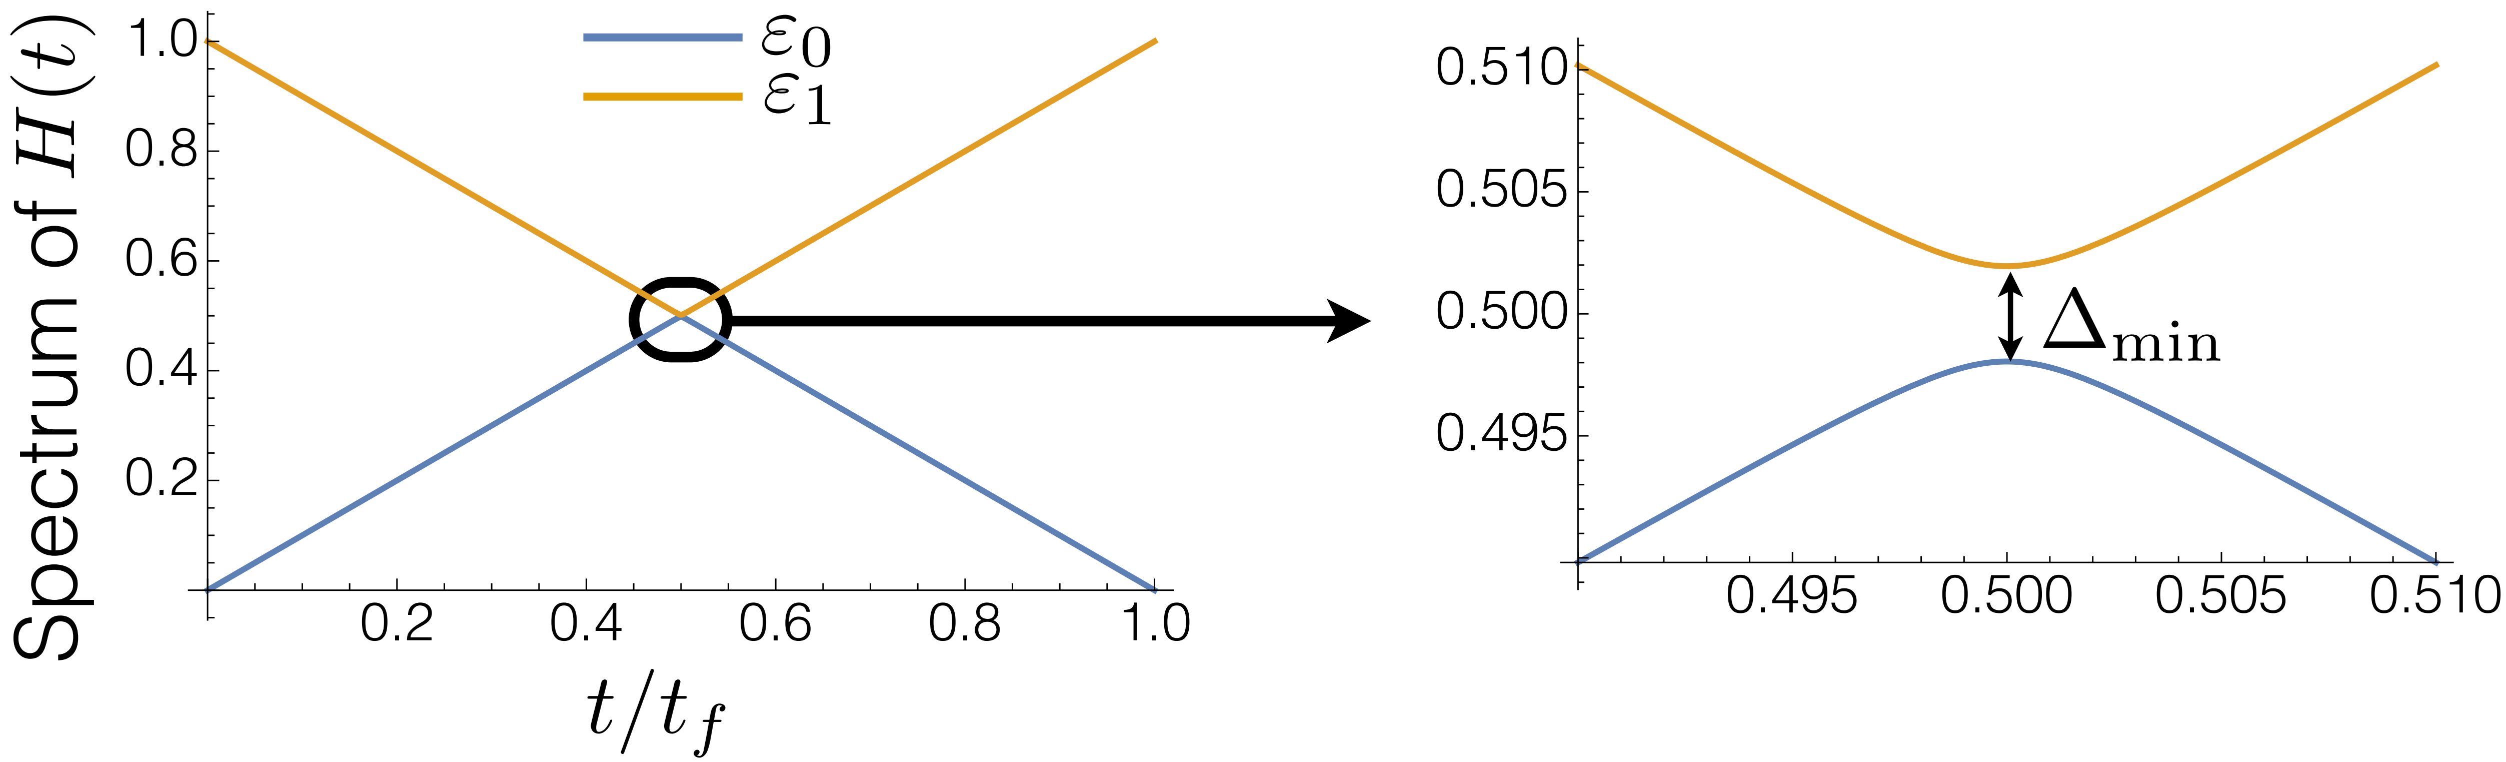

Adiabatic condition[1]

$$t_f \gg \frac{1}{\Delta_{\min}^3}$$

The efficiency of the algorithm is then determined by how t_f must scale with system size N

(1) S. Jansen,et al., J. Math. Phys. 48 (10), 102111 (2007)

An Alternative Picture

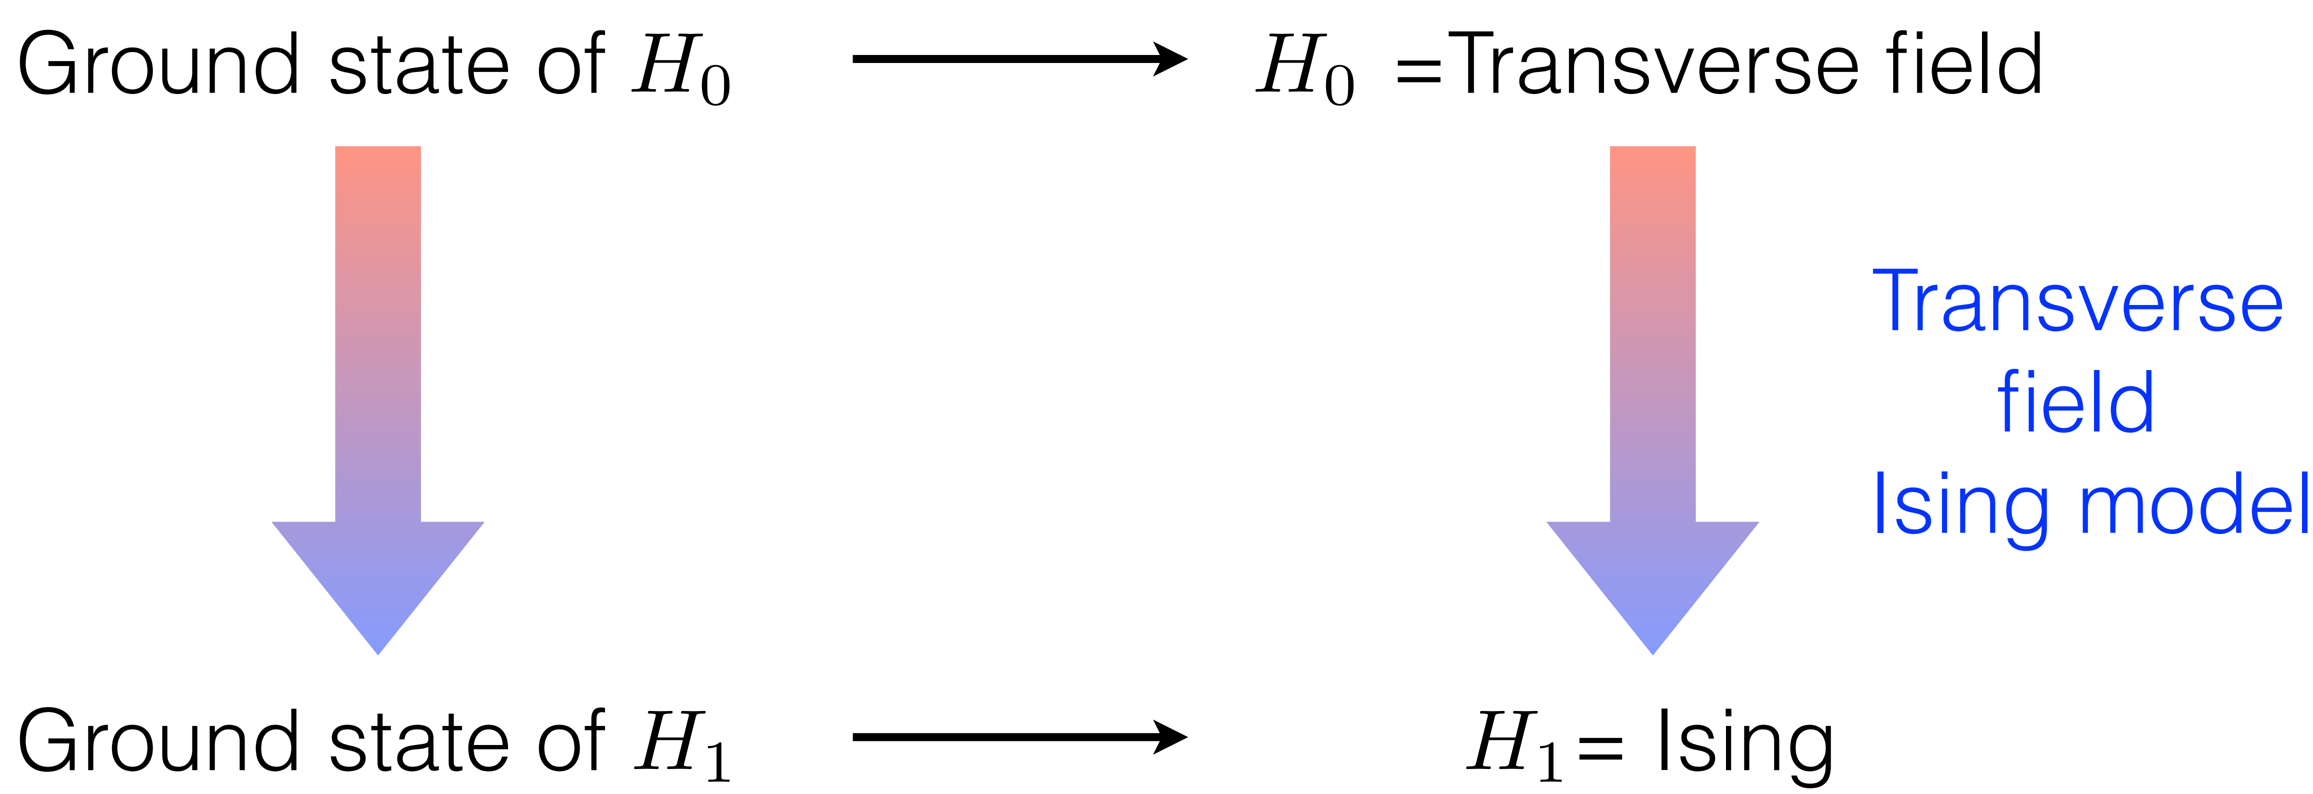

Phase diagram of a transverse field Ising model

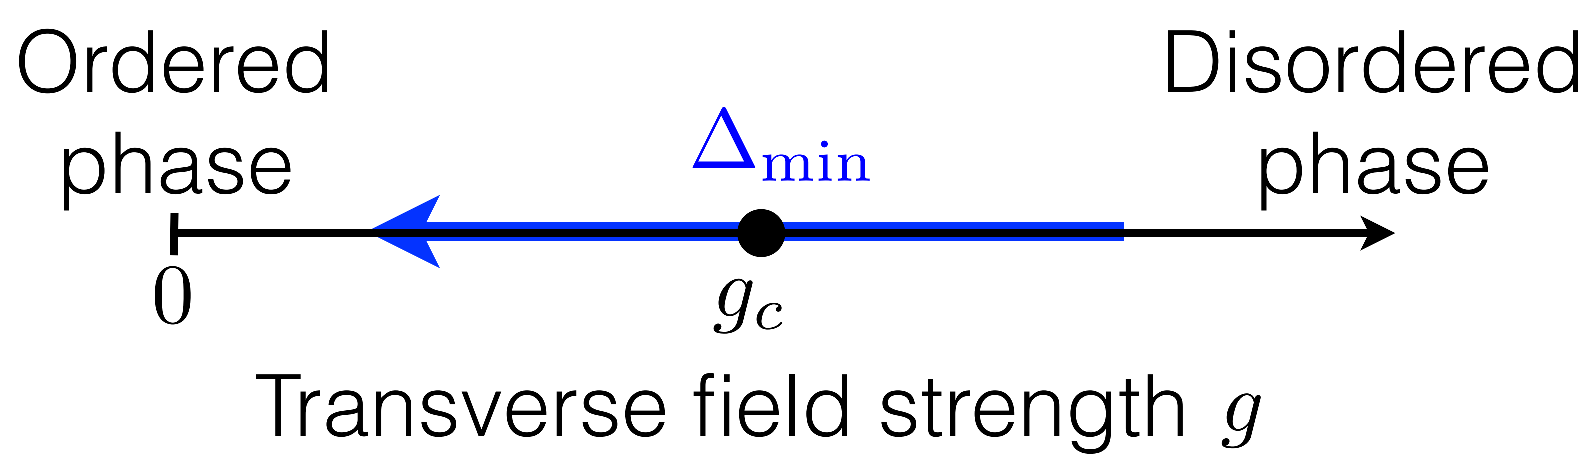

Quantum annealing tries to follow the ground state from the disordered phase to the ordered phase

A Classical Analogue

Quantum annealing (QA)[2]		Simulated annealing (SA)[1]
Ground state of H_0		High temperature state
Quench quantum fluctuations	Quantum/classical analogues	Quench thermal fluctuations
Ground state of H_1		Low temperature state

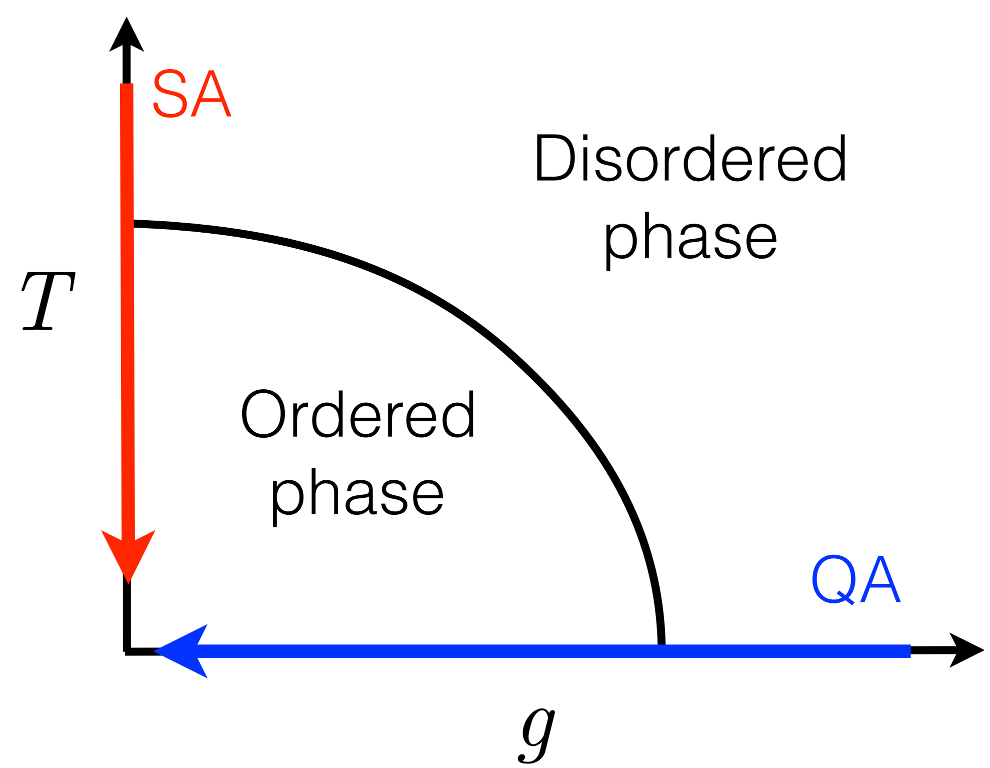

(1) S. Kirkpatrick et al., Science 220, 671–680 (1983)

(2) T. Kadowaki, and H. Nishimori, Phys. Rev. E 58 (5), 5355 (1998)

Simulated Annealing

High temperature state

Quench thermal fluctuations

Low temperature state

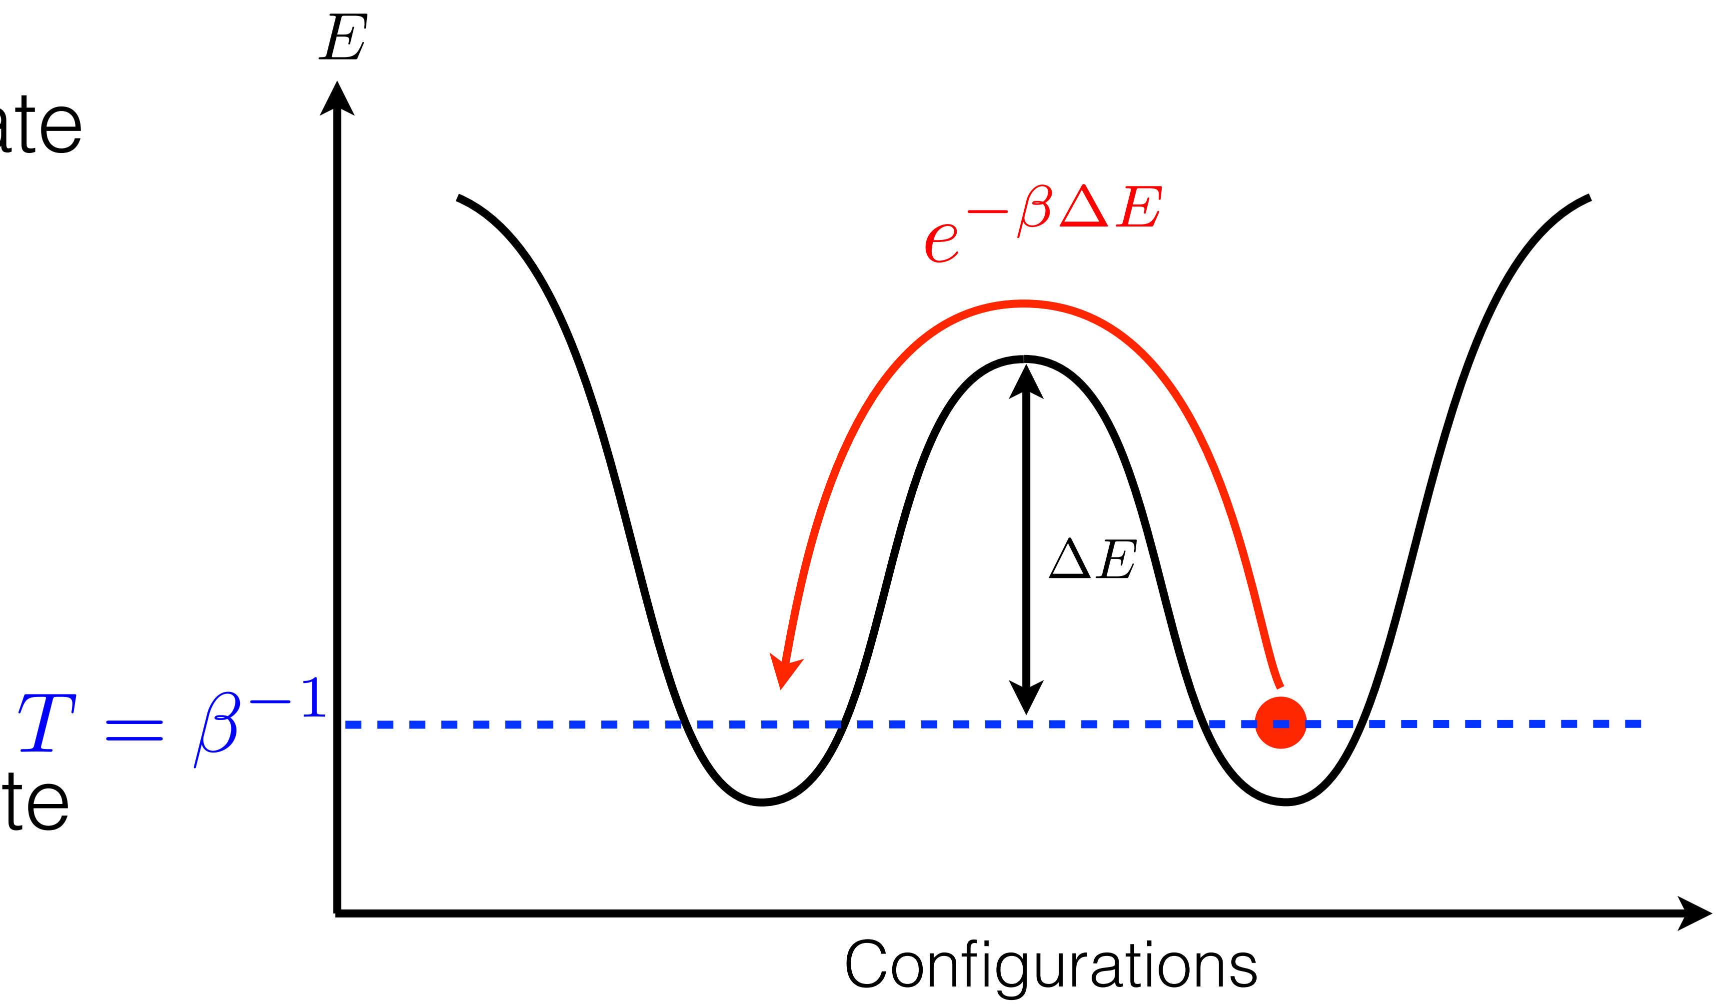

Thermally hopping over a barrier at low temperatures is exponentially suppressed in the barrier height ΔE

Why Might QA Have an Advantage over SA

Semiclassical potential evolves with time
State tries to stay in its local minimum

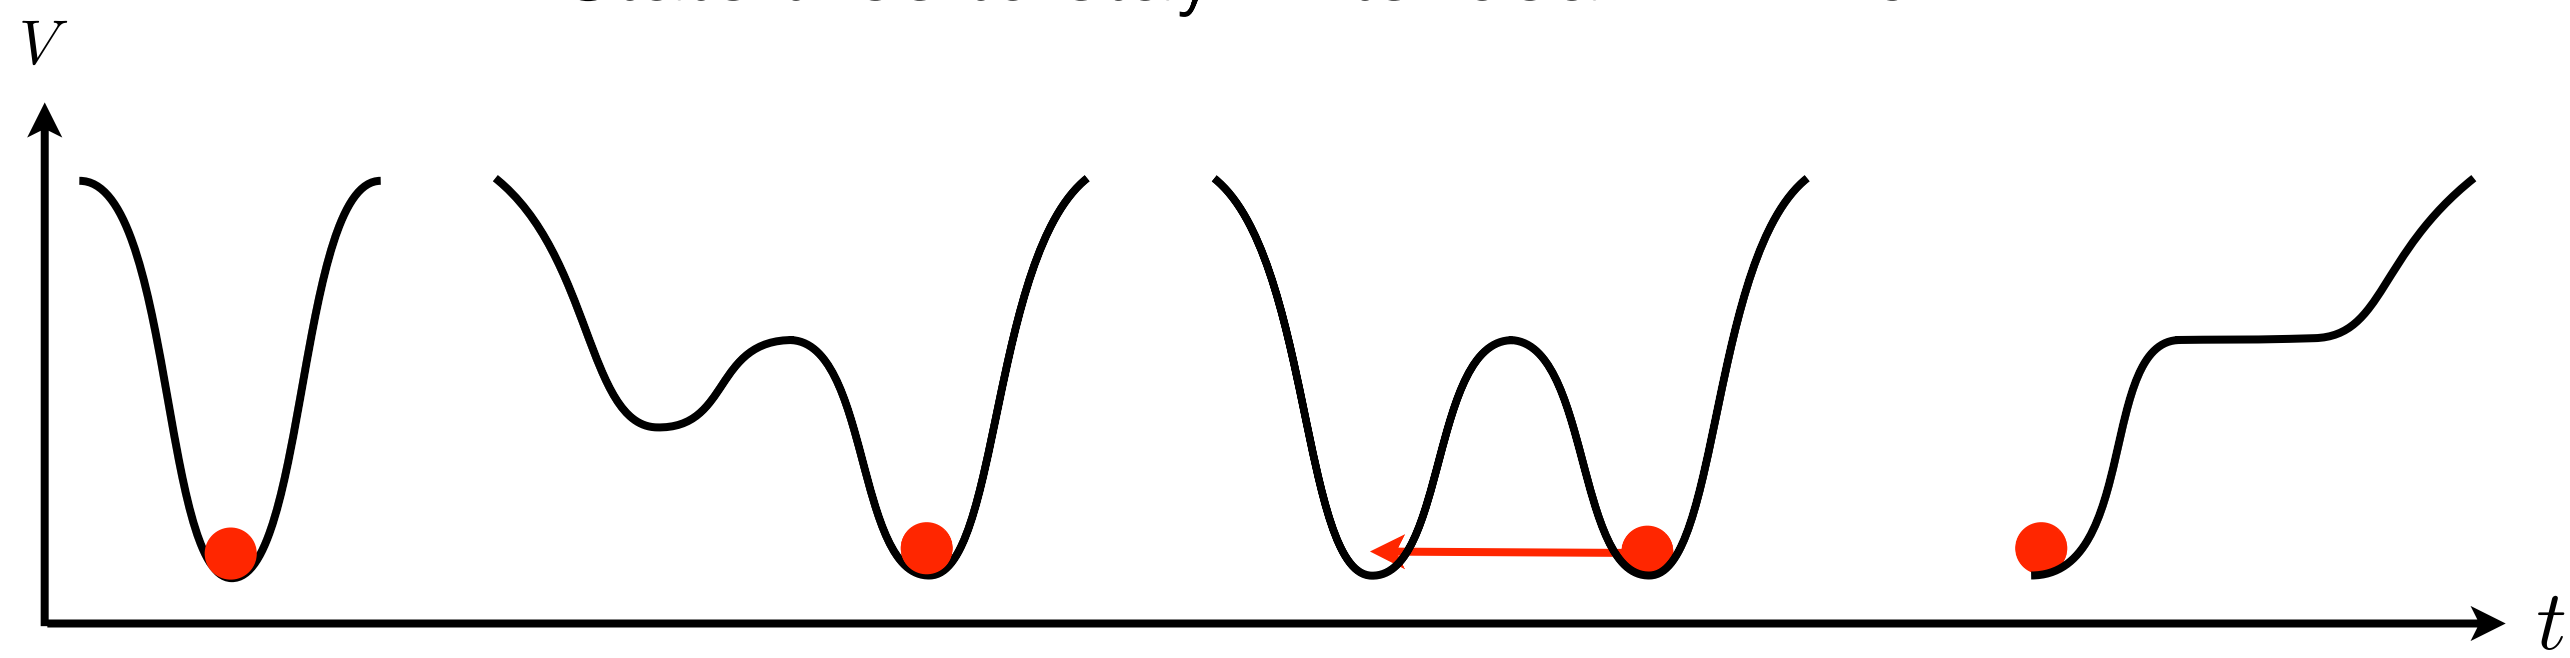

Small energy gaps are associated with double-well energy barriers through which the system must tunnel

The semiclassical potential may have a double-well, even if the classical potential does not[1]

Efficiency to traverse the barrier is related to how the barrier width and height scale with problem size

(1) S. Muthukrishnan, TA, and D. A. Lidar, Phys. Rev. X 6, 031010, (2016).

Illustrative Examples

Hamming Weight Problems

A classical Hamiltonian H_1 where the energy of a classical state only depends on the Hamming weight of that state (equivalent to counting the number of down-pointing spins)

Classical state	Hamming weight	Energy
$\|\uparrow\uparrow\uparrow\rangle$	0	E_0
$\|\downarrow\uparrow\uparrow\rangle$, $\|\uparrow\downarrow\uparrow\rangle$, $\|\uparrow\uparrow\downarrow\rangle$	1	E_1
$\vdots$	$\vdots$	$\vdots$
$\|\downarrow\downarrow\downarrow\rangle$	3	E_3

QA Hamiltonian with a transverse field is invariant under qubit permutations

Prototypical Hamming Weight Problem

The "spike" problem[1]

(shown here for $N = 16$)

Barrier width constant with problem size N

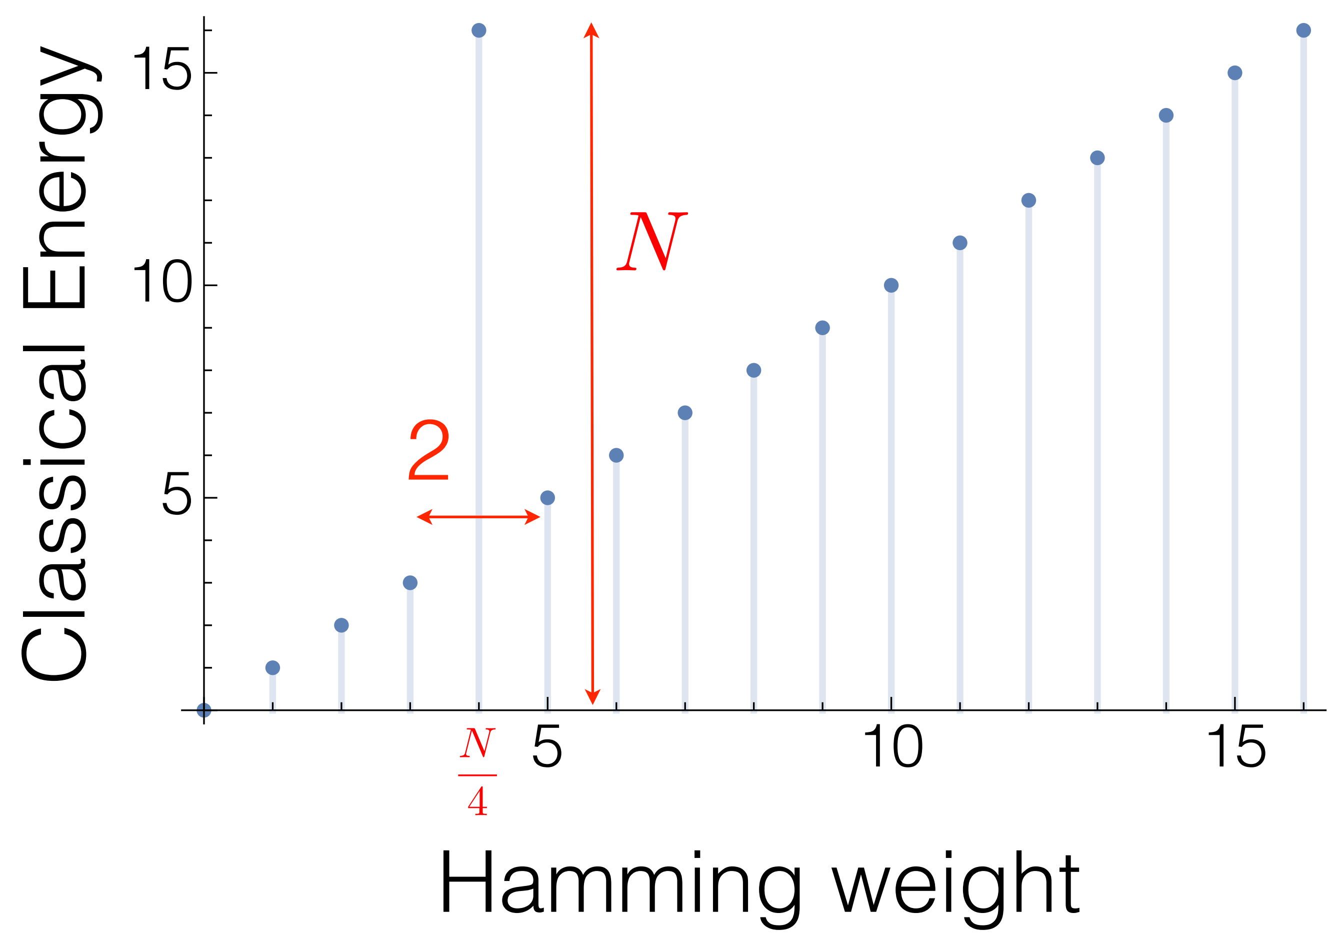

Barrier height grows with the problem size N

(single-spin) SA requires exponential time in the system size to find the ground state

(1) E. Farhi, J. Goldstone, and S. Gutmann, arXiv:quant- ph/0201031 (2002).

The "Spike" Problem

Classical landscape (landscape at $s = 1$)

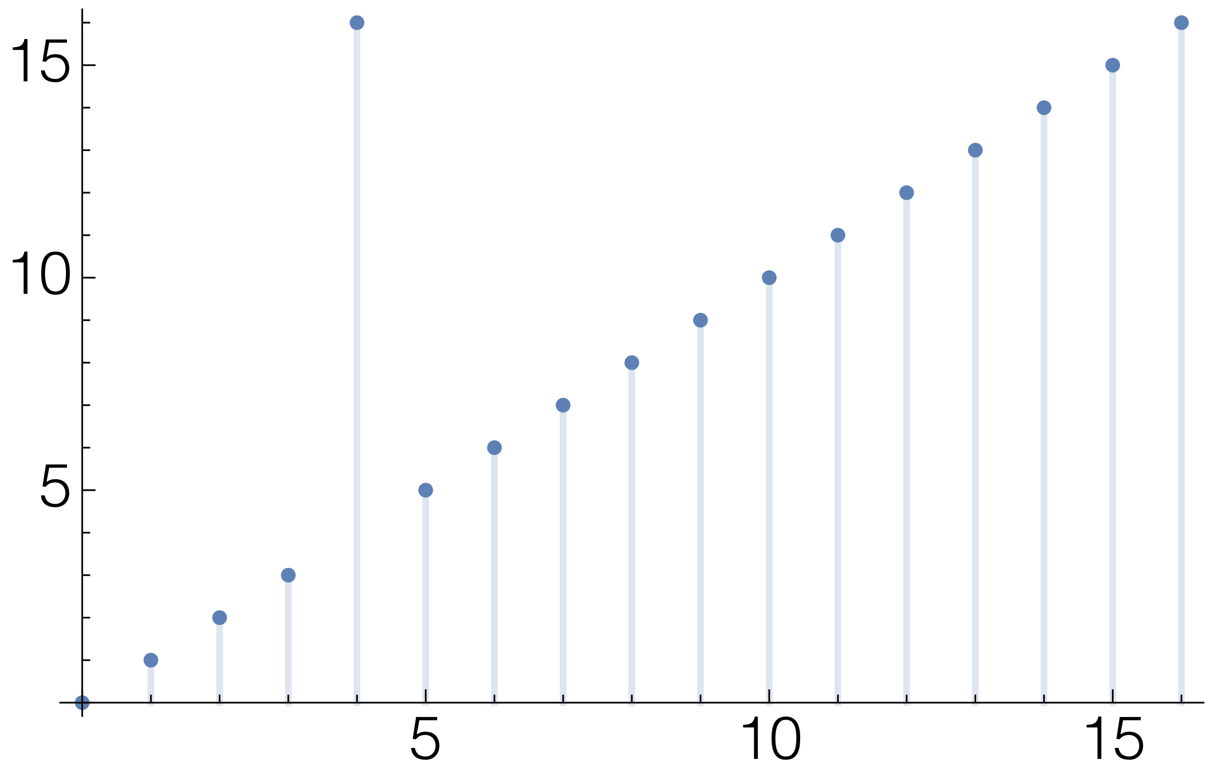

Quantum semi-classical landscape at $s = 0.359$

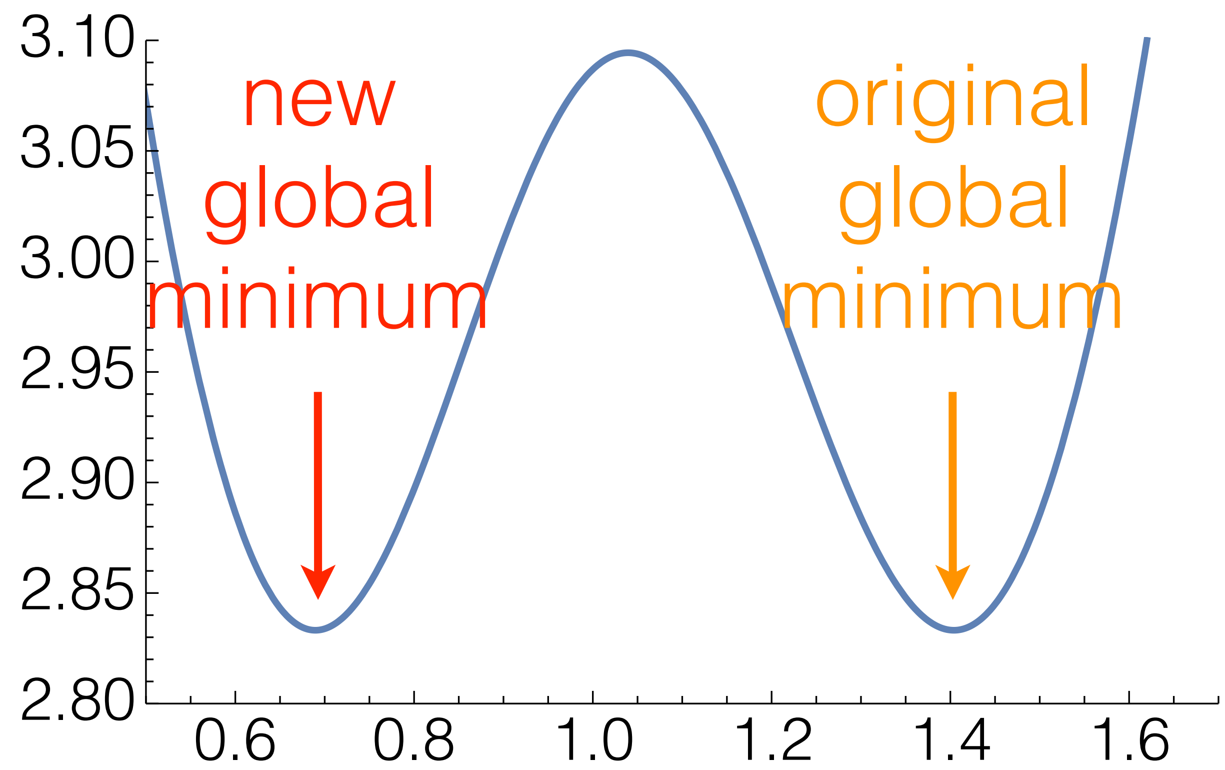

Double-well barrier height grows sub-linearly with N and width shrinks with increasing N

Quantum minimum gap scales polynomially with N

Exponential speedup over SA[1]

(1) E. Farhi, J. Goldstone, and S. Gutmann, arXiv:quant- ph/0201031 (2002).

Recap: Closed System Quantum Annealing

Quantum annealing

Prepare system in the ground state of H_0 → Adiabatic evolution guarantees success → Ground state of H_1 encodes the solution

Efficiency controlled by the scaling of the minimum gap

The changing energy landscape can be very different from the classical energy landscape

↓

Can lead to quantum advantage over classical algorithms like simulated (thermal) annealing

So What's the Problem

Hamming weight Hamiltonians examples require N-body operators.

The intuition gained has not yet led to the demonstration of a speedup for physically realizable problems

Why not?

- No symmetries to help facilitate the analysis
- Simulations are not possible at large sizes
- Very good classical algorithms

Only way to test for speedup may be to build a device and see (?)

The Realistic Setting

A physical device naturally couples the system to additional degrees of freedom (the “bath”)

Decoherence

~~Closed~~ Open system quantum annealing

↓

Unitary and dissipative dynamics

(non-unitary) population transfer

decay of phase relationship between specific states

Conflict between running adiabatically (long time) and minimizing the interaction time with the bath

Weakly Coupled Open Quantum Systems

The most innocuous model of decoherence for QA [(1)]
weak coupling and Markovian

Key feature

Dissipative dynamics pushes system towards the thermal state[(2)]

↓

$$\lim_{t_f \to \infty} \rho(t) \to \rho_G(t)$$ [(3)]

$$\rho_{\mathrm{G}}(t) = \frac{\exp\left(-\beta H(t)\right)}{\mathrm{Tr}\left[\exp\left(-\beta H(t)\right)\right]}$$

Affect on QA

Minimum energy gap decreases with increasing problem size, so adiabatic theorem requires us to use longer annealing times

More time spent interacting with the environment means system gets closer and closer to the thermal state

(1) TA and D. A. Lidar, Phys. Rev. A 91, 062320 (2015)
(2) TA, S. Boixo, D. A. Lidar, P. Zanardi, New J. Phys. 14, 123016 (2012)
(3) L. C. Venuti, TA, D. A. Lidar, and P. Zanardi Phys. Rev. A 93, 032118 (2016)

Thermal Quantum Annealer

For a thermalizing quantum annealer

Success depends on how much weight the thermal state has on the ground state

Distribution becomes more gaussian-like, with a mean that moves away from the origin

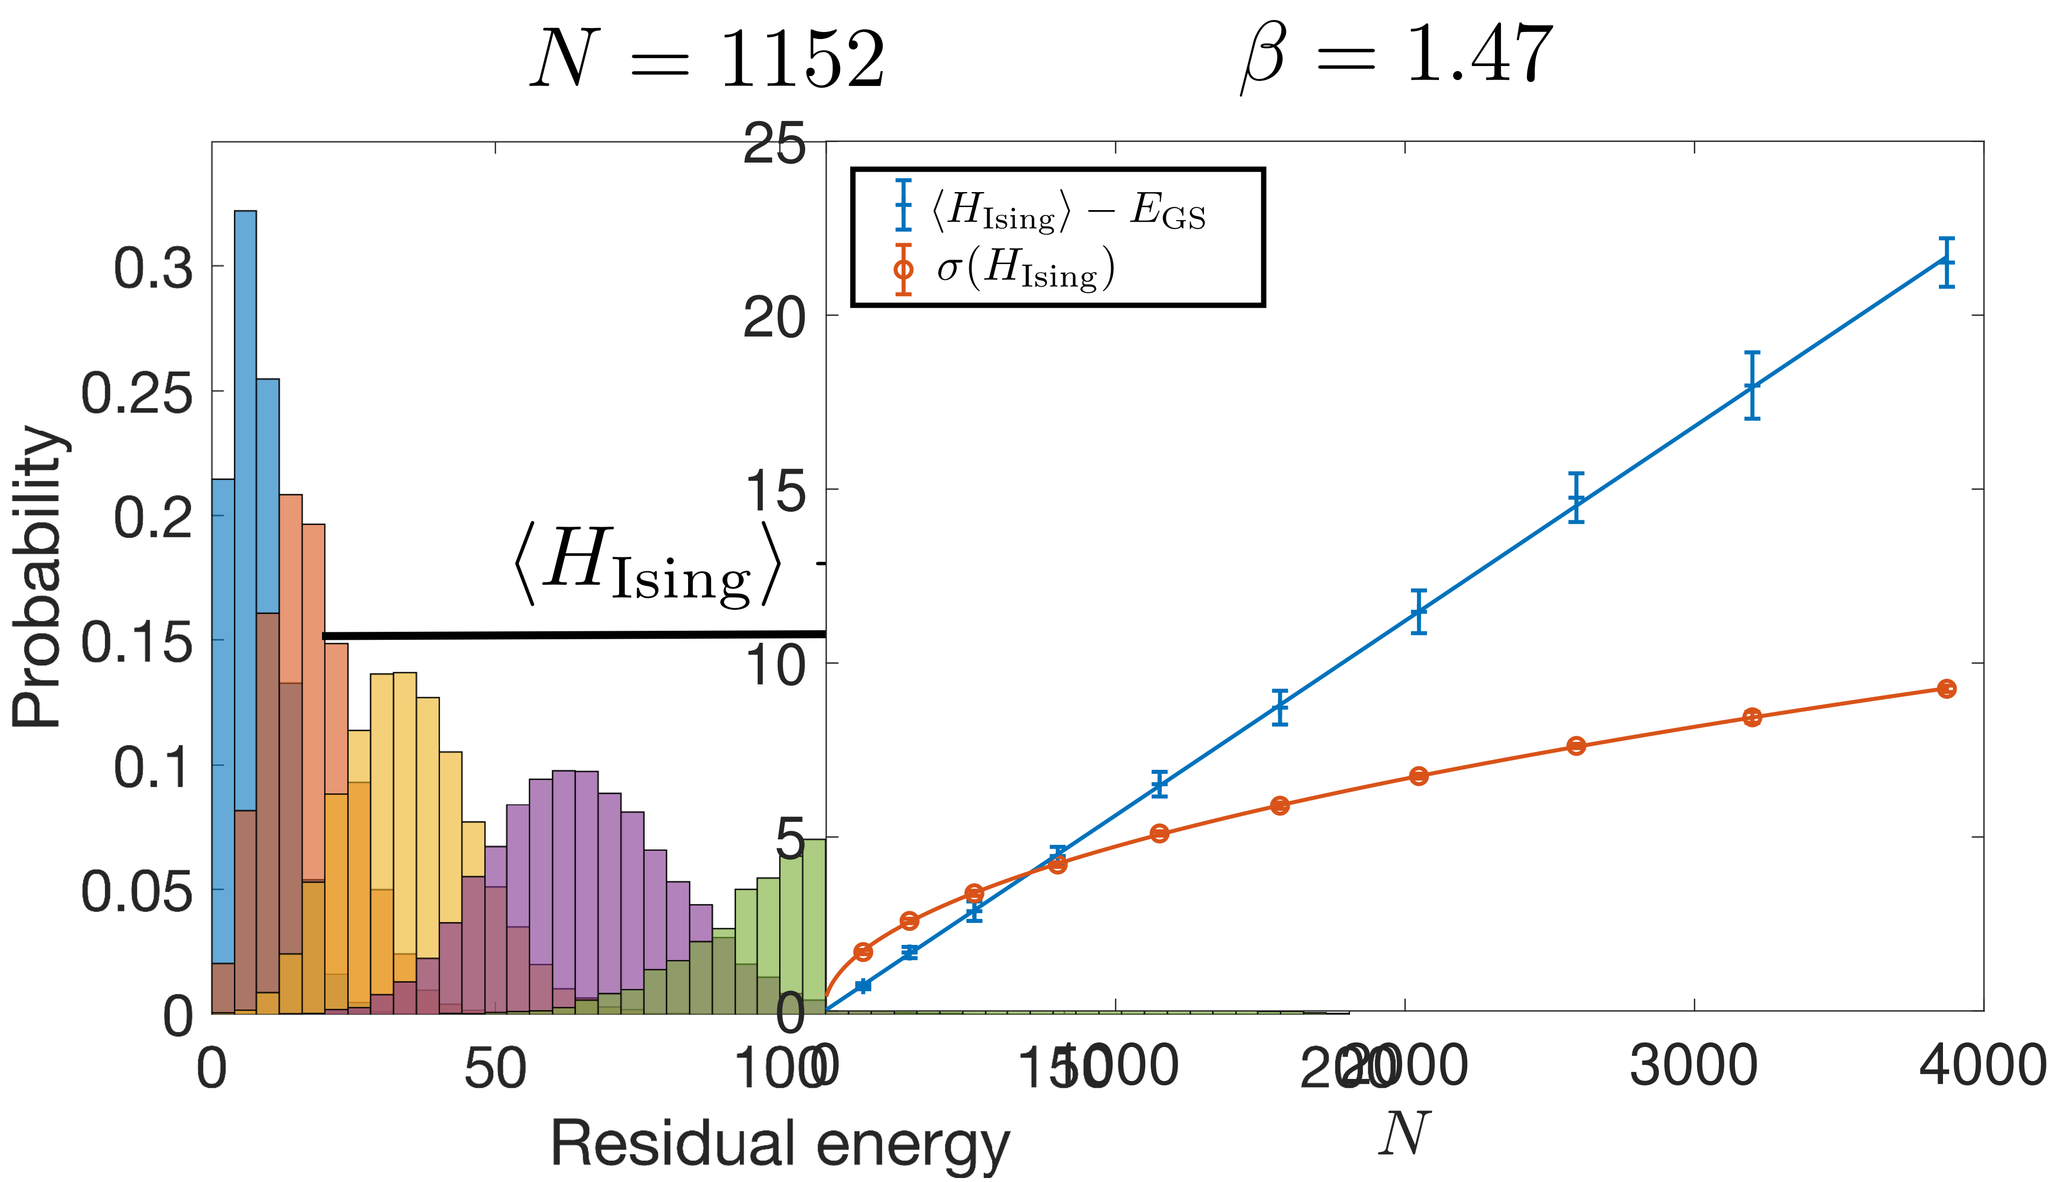

At a fixed temperature, it becomes exponentially unlikely to sample the ground state with increasing problem size(1)

(1) TA, V. Martin-Mayor, I. Hen, Phys. Rev. Lett. 119, 110502 (2017).

Analog Devices leads to Analog Errors

Analog control of the Hamiltonian leads to misspecification of the final Hamiltonian

Desired

$$H_{\text{Ising}} = \sum_i h_i \sigma_i^z + \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z$$

Implemented

$$H'_{\text{Ising}} = \sum_i (h_i + \delta h_i) \sigma_i^z + \sum_{\langle i,j \rangle} (J_{ij} + \delta J_{ij}) \sigma_i^z \sigma_j^z$$

If the perturbation is sufficiently large, the ground state H'_{Ising} is different from the ground state of H_{Ising}

Even a perfect adiabatic evolution will result in the wrong answer!

Implementation Errors

How does a fixed precision affect the success of QA with increasing problem size?

Define p_{S} to be the probability that the ground state of H'_{Ising} matches one of the ground states of H_{Ising}

Pick a noise model

$$\delta h_i, \delta J_{ij} \sim \mathcal{N}(0, \sigma^2)$$

Use instances at different sizes N with known ground states, and generate many noisy realizations of the same instances

Scaling with Implementation Errors

Instances defined with a Chimera connectivity with range 3 and only two ground states[1]

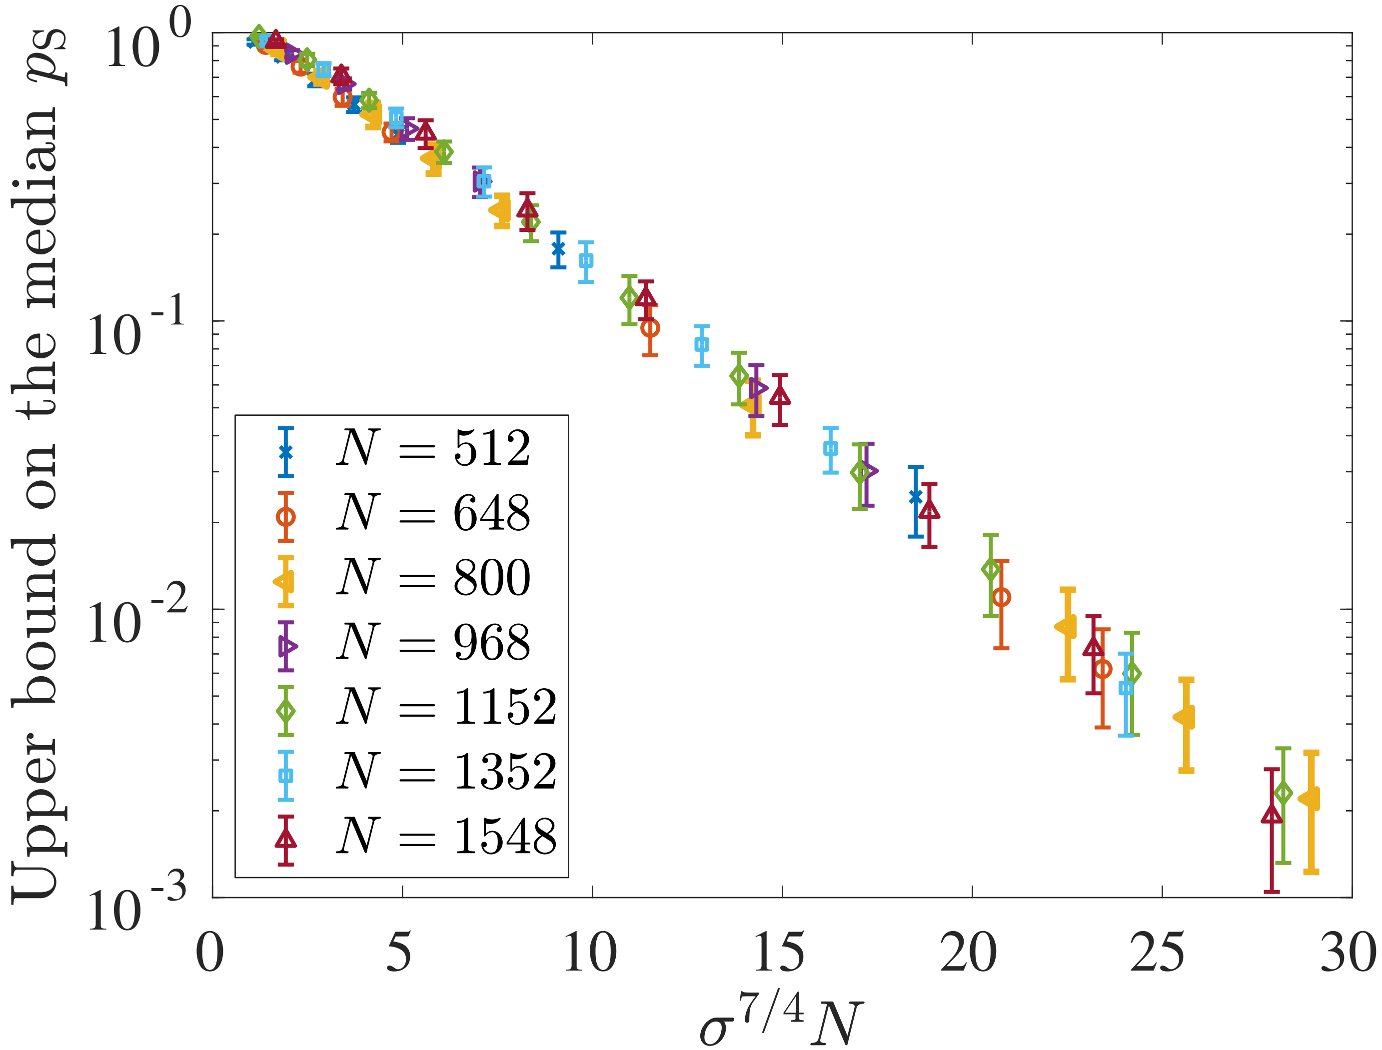

For a fixed noise strength σ, the probability of the ground state not changing decreases exponentially with problem size N

Different problem classes exhibit different dependence on σ

(1) TA, V. Martin-Mayor, I. Hen, arXiv:1806.03744.

Recap: Open System Quantum Annealing

Even in the optimistic setting of decoherence

Temperature must be scaled down with increasing problem size for thermalizing quantum annealer

$$T_{\mathrm{S}}^{(\text{worst case})} \sim \frac{1}{N^{\alpha}}$$

Implementation errors must be scaled down with increasing problem size, even for an otherwise perfect device

$$\sigma_{\mathrm{S}}^{(\text{worst case})} \sim \frac{1}{N}$$

Scalable quantum annealing has no hope without fault tolerant quantum error correction, which remains on open theoretical question

Commercially Available Quantum Annealers

D-Wave Systems has been selling (purported) quantum annealing processors since 2011

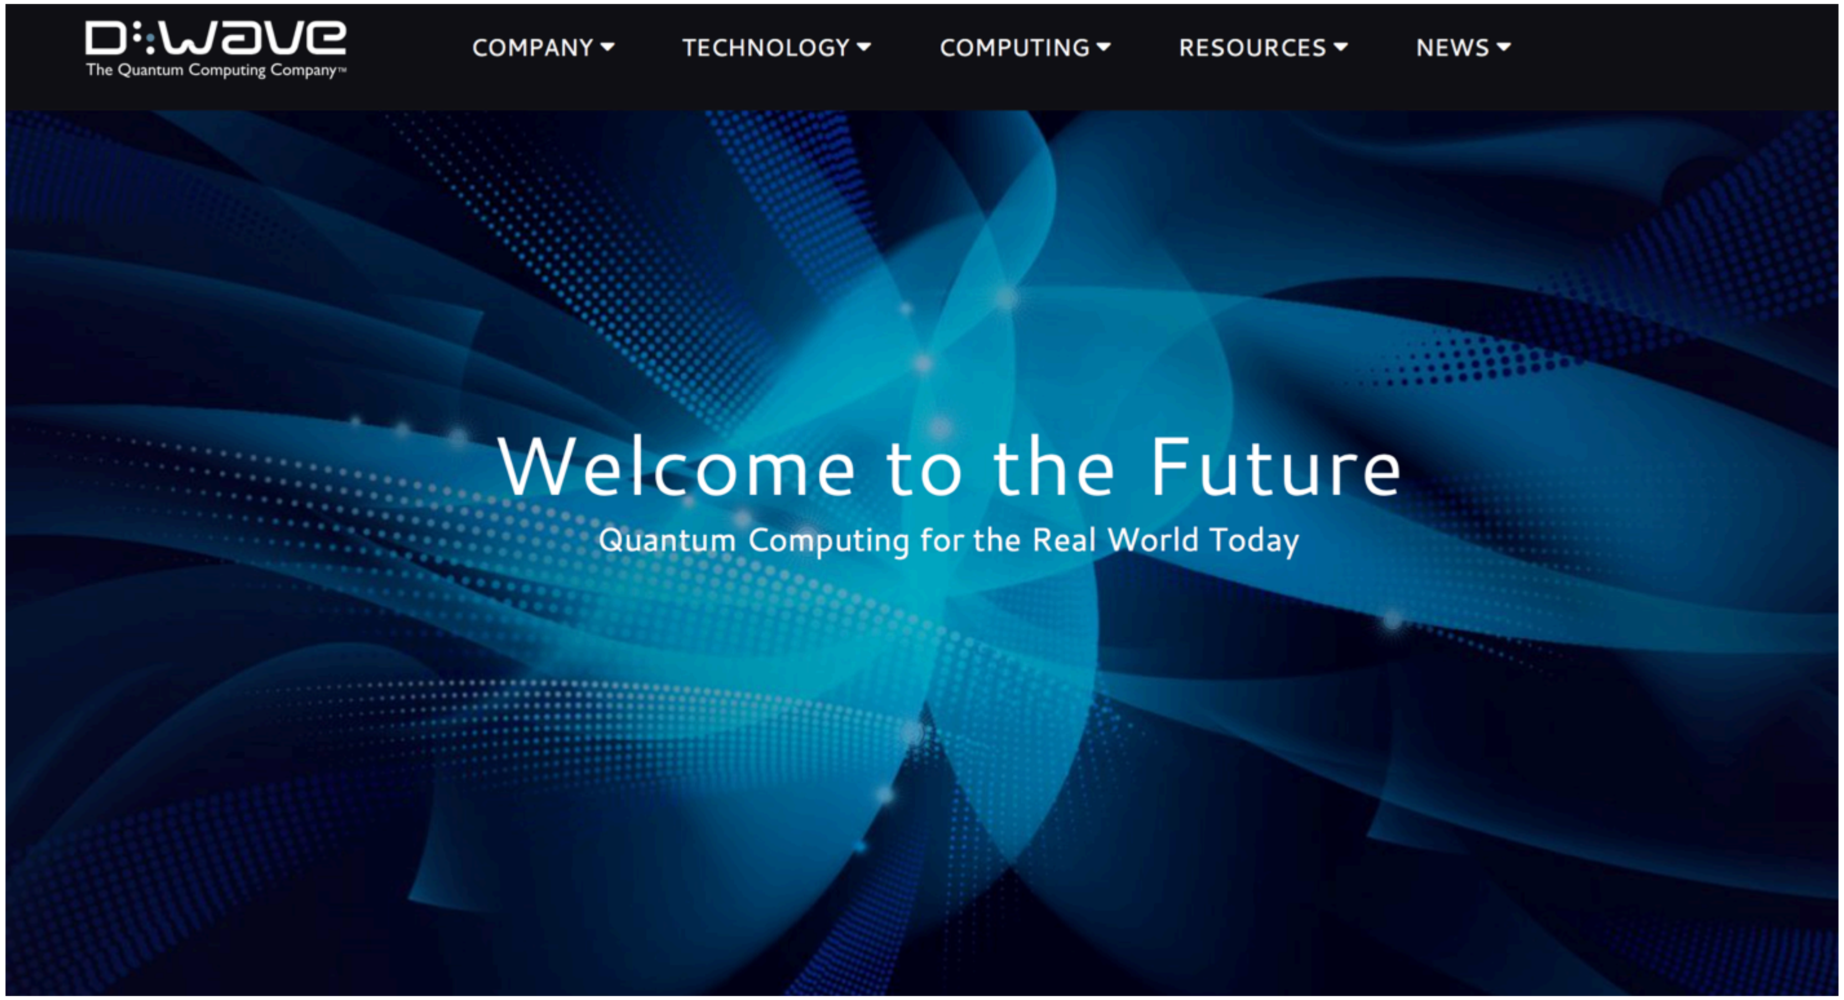

2011: D-Wave 1 “Rainier”	128 qubits	20mK
2013: D-Wave 2 “Vesuvius”	512 qubits	17mK
2015: D-Wave 2X “Washington”	1152 qubits	13mK
2017: D-Wave 2000Q	2048 qubits	12-15mK

How the D-Wave Processor Works

End user programs the Ising local fields $\{h_i\}$
and Ising couplers $\{J_{ij}\}$

$$H_{\text{Ising}} = \sum_i h_i \sigma_i^z + \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z$$

Ising connectivity limited to the physical qubit connectivity of the device

Square grid of 8 qubit unit cells

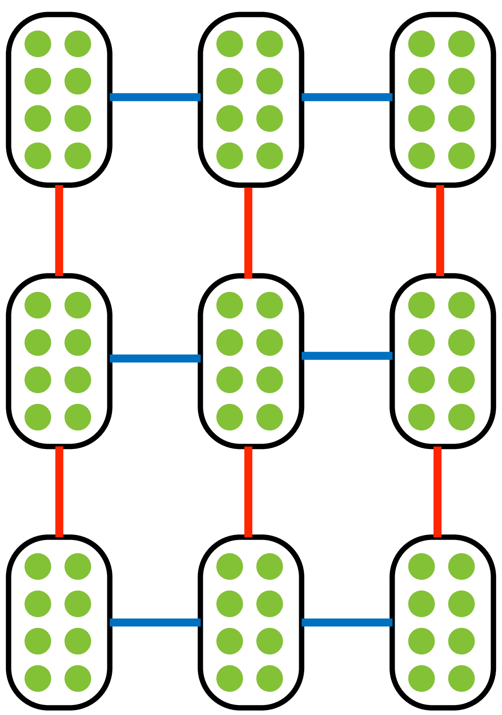

Unit cell

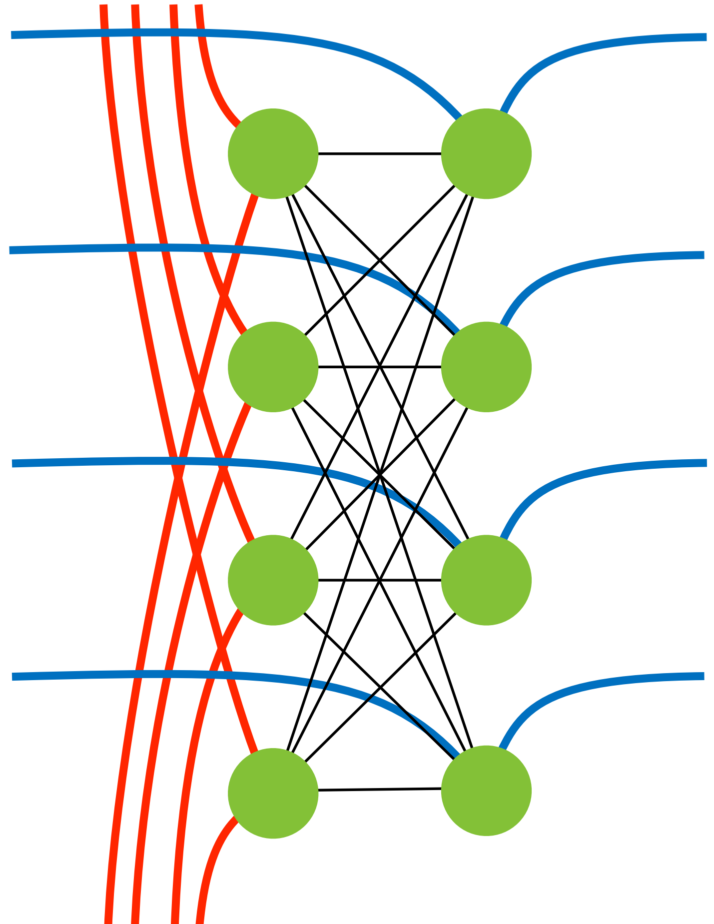

Benchmarking D-Wave Processors

Several benchmarking studies over several generations of D-Wave devices defined on $L \times L$ Chimera graphs ranging from 128 to 2048 qubits

DW1, $L = 4$ (128 qubits)	DW2, $L = 8$ (512 qubits)	DW2X, $L = 12$ (1152 qubits)	DW2000Q, $L = 16$ (2048 qubits)
Boixo *et al*. (2014)	Rønnow *et al*. (2014) Hen *et al*. (2015) King *et al*. (2015)	King *et al*. (2015) Denchev *et al*. (2016)	King *et al*. (2017) TA *et al*. (2018) Mandrà *et al*. (2018)

Benchmarking standards set in Rønnow *et al*. (2014)

Computational cost measured in terms of time-to-solution (TTS), required time to run algorithm to find the ground state at least once with a 0.99 probability

At each problem size, algorithm parameters are optimized to minimize the TTS (averaged over instances), $\langle \mathrm{TTS} \rangle^*$

Has it demonstrated a speedup?

No demonstration of a scaling advantage over classical algorithms(1)

DW2KQ: D-Wave 2000Q device

SA: Simulated Annealing

PIQMC: Path-Integral Monte Carlo Annealing (aka Simulated quantum annealing)

PT-ICM: Parallel tempering with isoenergetic cluster moves

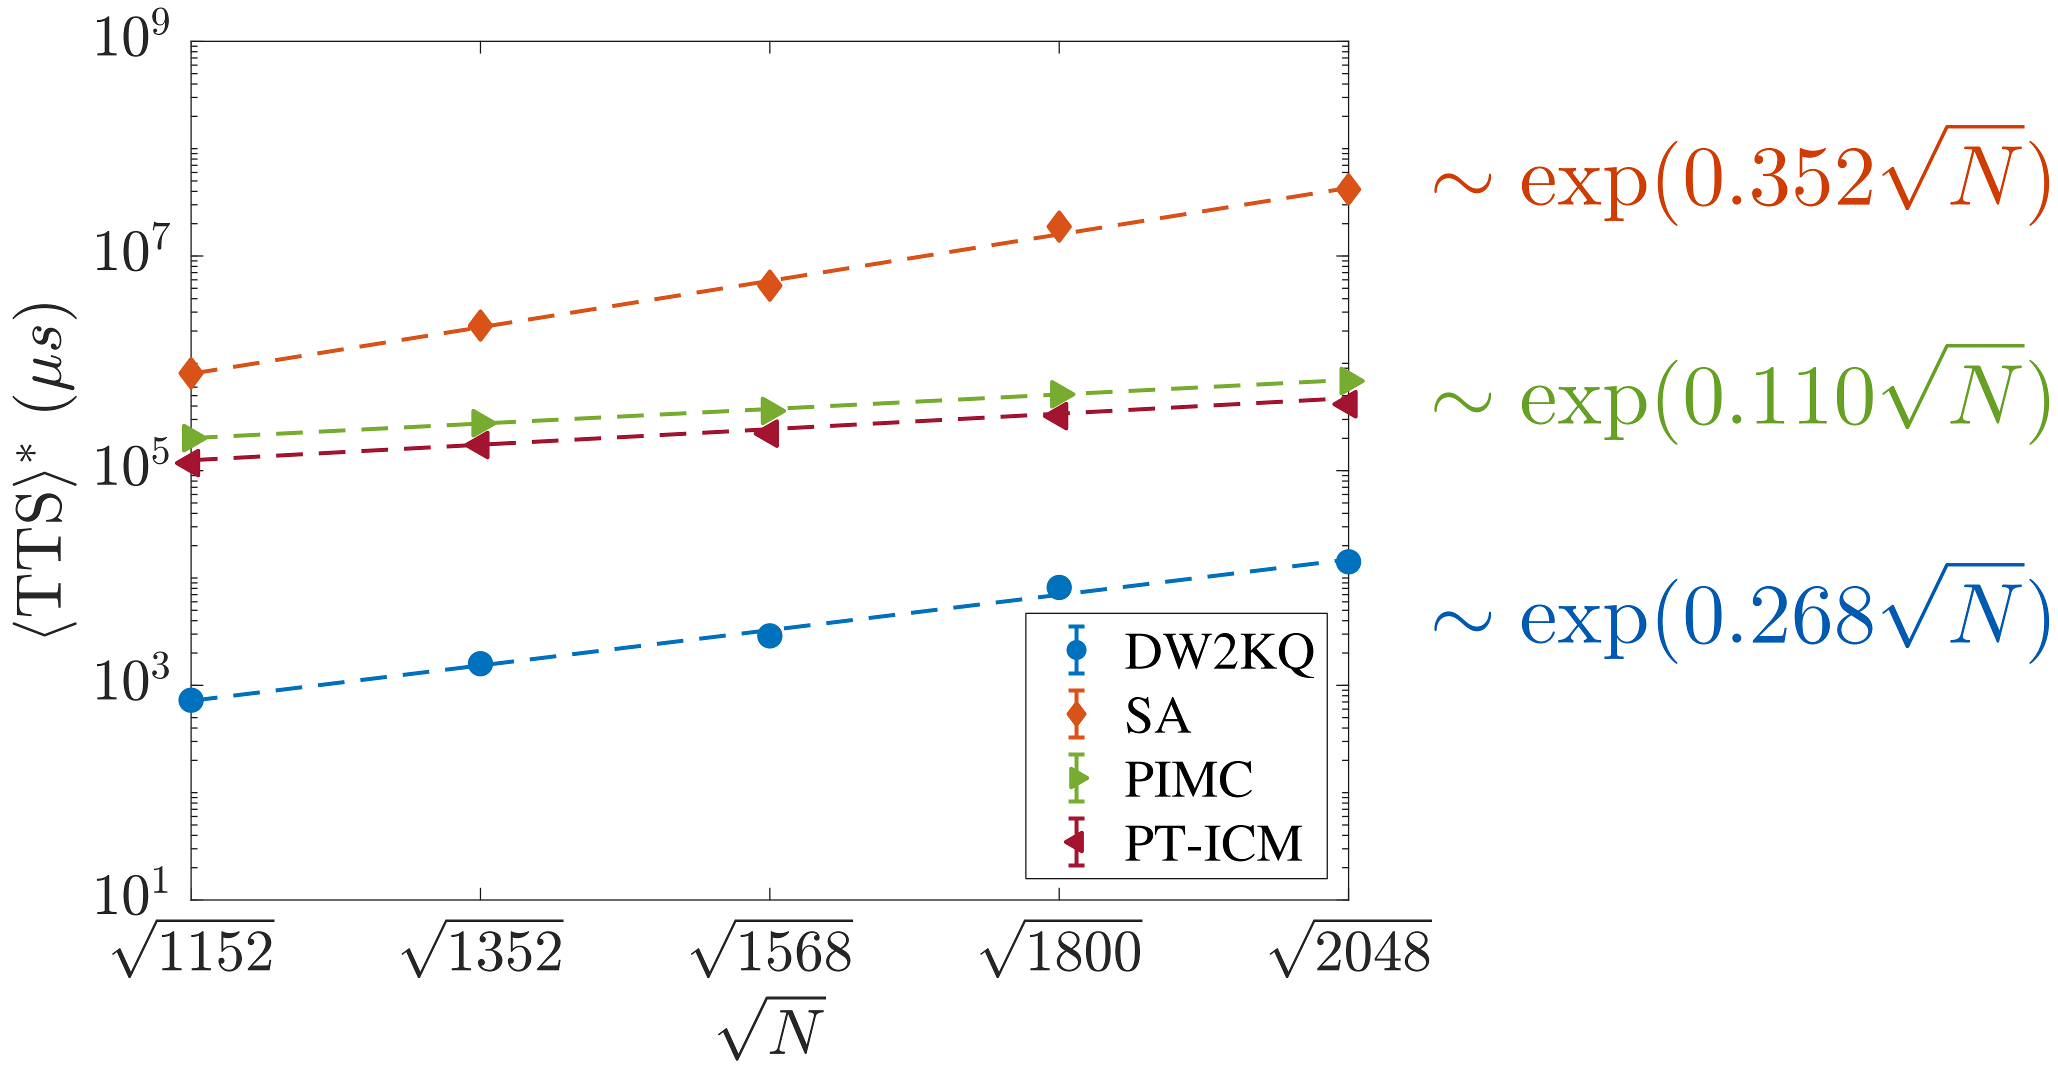

Classical algorithms can be very efficient at solving this class of Ising problems on the D-Wave connectivity graph(1)

(1) TA and D. A. Lidar, Phys. Rev. X 8, 031016 (2018).

Temperature slow down?

The device temperature is too high

Consider the scaling of path-integral Monte Carlo annealing at two different temperatures using the D-Wave annealing schedule

D-Wave operating temperature corresponds to $\beta = 0.51$. D-Wave scaling is $\exp(0.268\sqrt{N})$

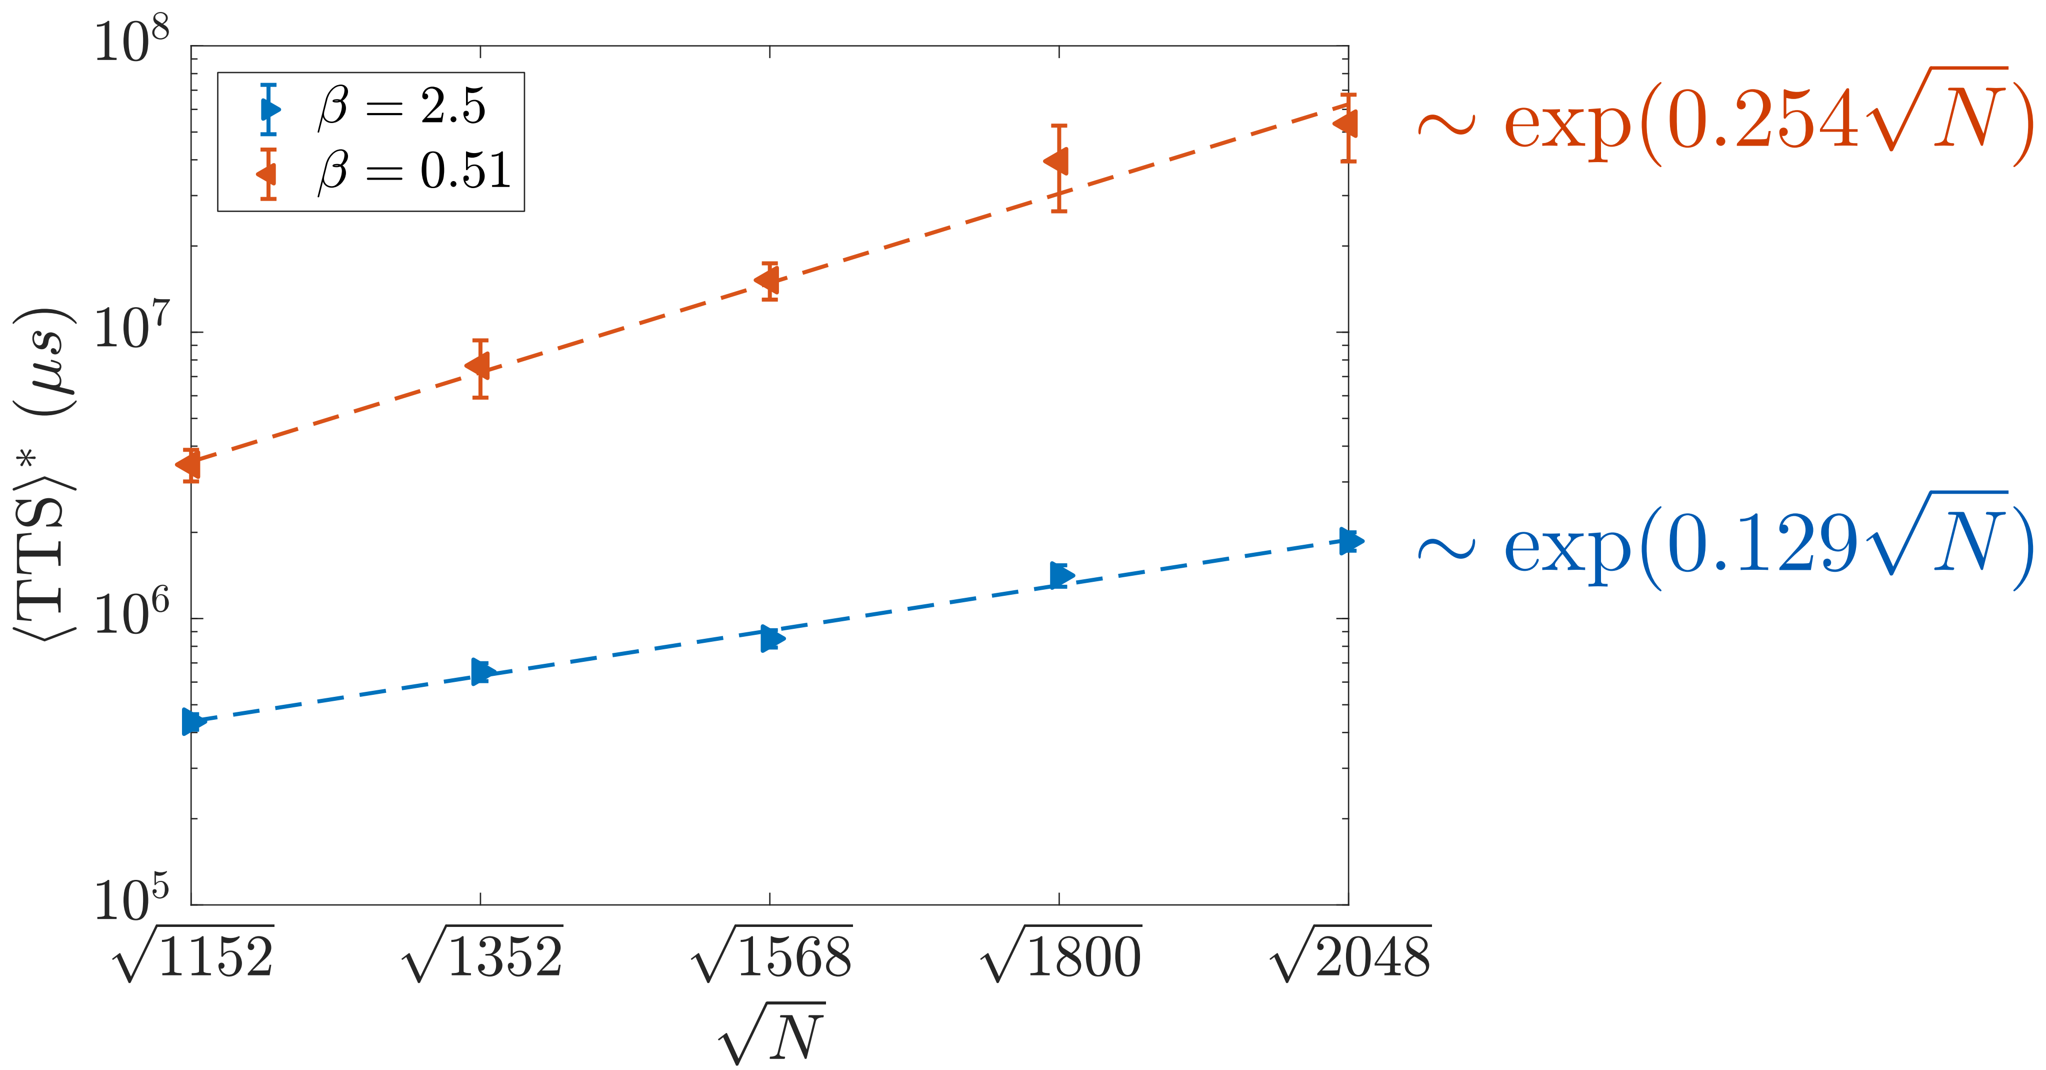

Performance improves dramatically at lower temperatures

Prefactor versus Scaling

I have so far focused on the scaling with problem size

Even if we don't get a scaling advantage, we may get a wall-clock time advantage (prefactor advantage)[1,2]
Requires accounting of all time costs (initial state preparation, measurement time, etc.)

Still no clear evidence of this advantage yet on current devices

Other possible metrics:
Power consumption per quality of solution[3]

(1) S. V. Isakov et al. PRL 117, 180402 (2016).
(2) Z. Jiang et al. PRA 95, 012322– (2017).
(3) S. Mandrà and H. Katzgraber, Quant. Sci. Technol. 3, 04LT01 (2018)

'Standard' Quantum Annealing

Discussion so far has been about using transverse field Hamiltonian to drive the anneal

$$H(s) = (1-s)\left(-\sum_i \sigma_i^x\right) + sH_{\text{Ising}}$$

To date,
no theoretical or experimental evidence for a quantum speedup for Ising-type Hamiltonians

Nothing stops us from going beyond the standard setup

Beyond Standard QA

Introduce intermediate 'catalyst' Hamiltonian H_c to help the anneal[1]

$$H(s) = (1-s)H_0 + s(1-s)H_\mathrm{c} + sH_\mathrm{Ising}$$

$$H_\mathrm{c} = \alpha \sum_{\langle i,j \rangle} \sigma_i^x \sigma_j^x$$

$\alpha = -1$

Ferromagnetic transverse interaction

Hamiltonian $H(s)$ remains stoquastic[2]

$\alpha = 1$

Anti-ferromagnetic transverse interaction

Hamiltonian $H(s)$ may be non-stoquastic[2]

In most cases, $\alpha = -1$ is the better choice, and often $\alpha = 1$ is worse than $\alpha = 0$[1,3]

(1) E. Crosson, et al., arXiv preprint arXiv:1401.7320 .
(2) S. Bravyi, et al., Quant. Inf. Comp. 8, 0361 (2008).
(3) L. Hormozi, et al., Phys. Rev. B 95, 184416 (2017).

Non-generic Counterexample

Infinite-range ferromagnetic p-spin model[1]

$$H(s,\lambda) = -(1-s)\sum_i \sigma_i^x + s(1-\lambda)n^{-1}\left(\sum_i \sigma_i^x\right)^2 - s\lambda n^{1-p}\left(\sum_i \sigma_i^z\right)^p$$

$p = 6$[2]

$\alpha = -1$

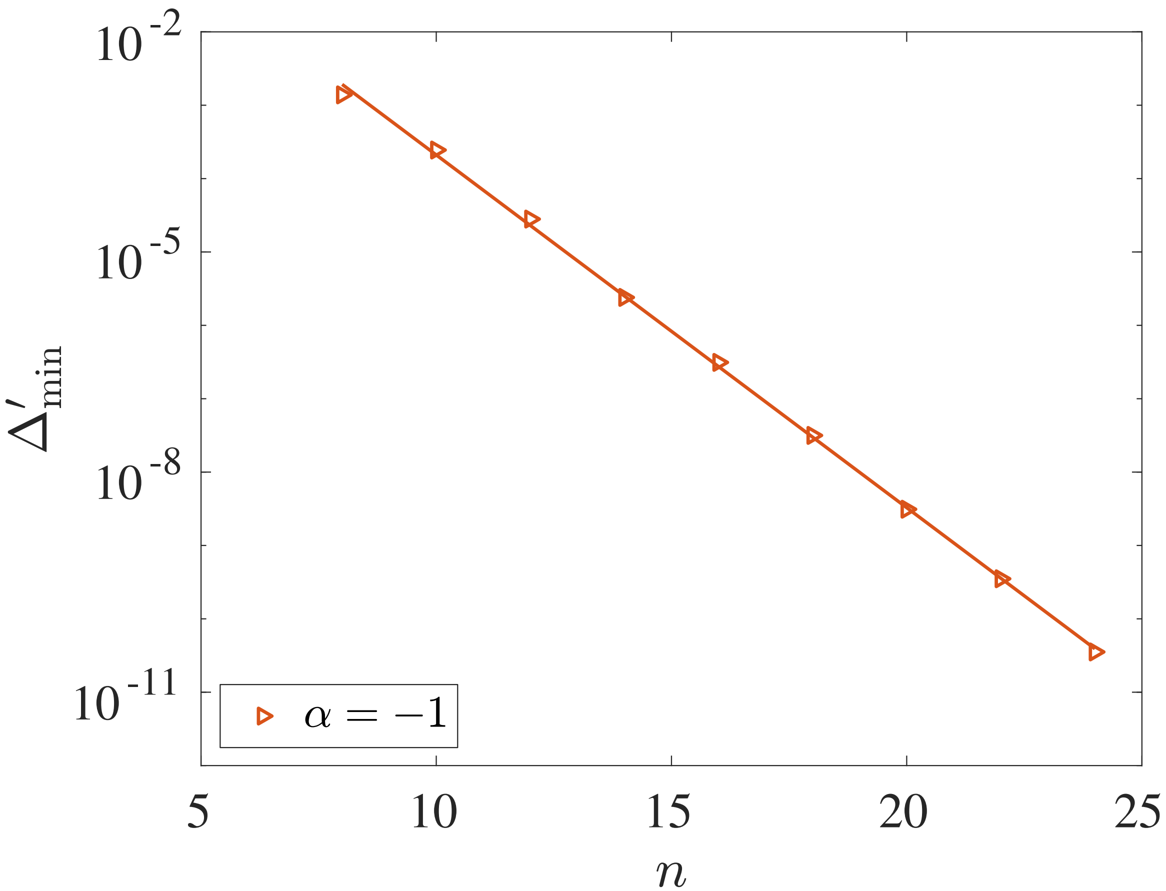

Exponentially closing gap

$\alpha = 1$

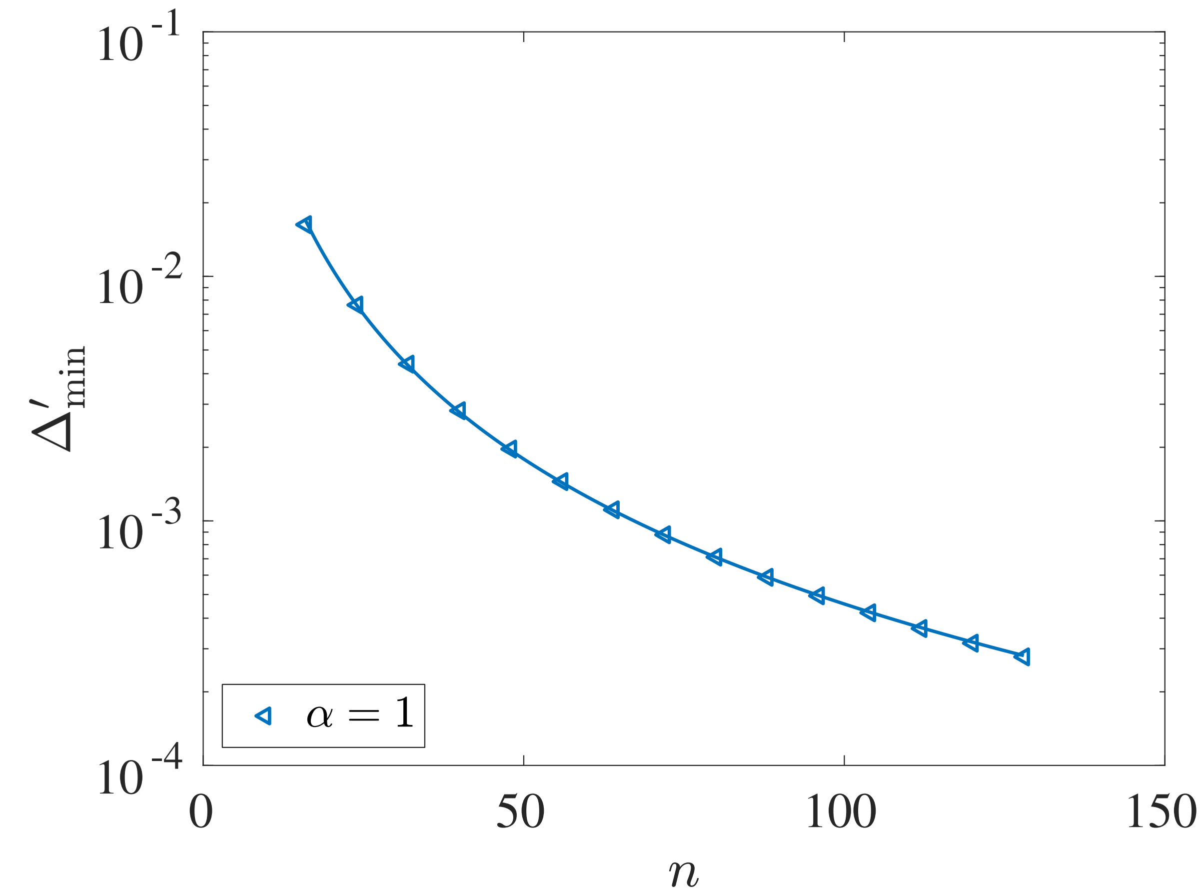

Polynomially closing gap

(1) Y Seki and H Nishimori, Phys. Rev. E 85, 051112 (2012).
(2) TA, arXiv:1811.09980.

Incremental steps vs One step

Ground state changes incrementally along the anneal for $\alpha = 1$, as opposed to suddenly for $\alpha = -1$

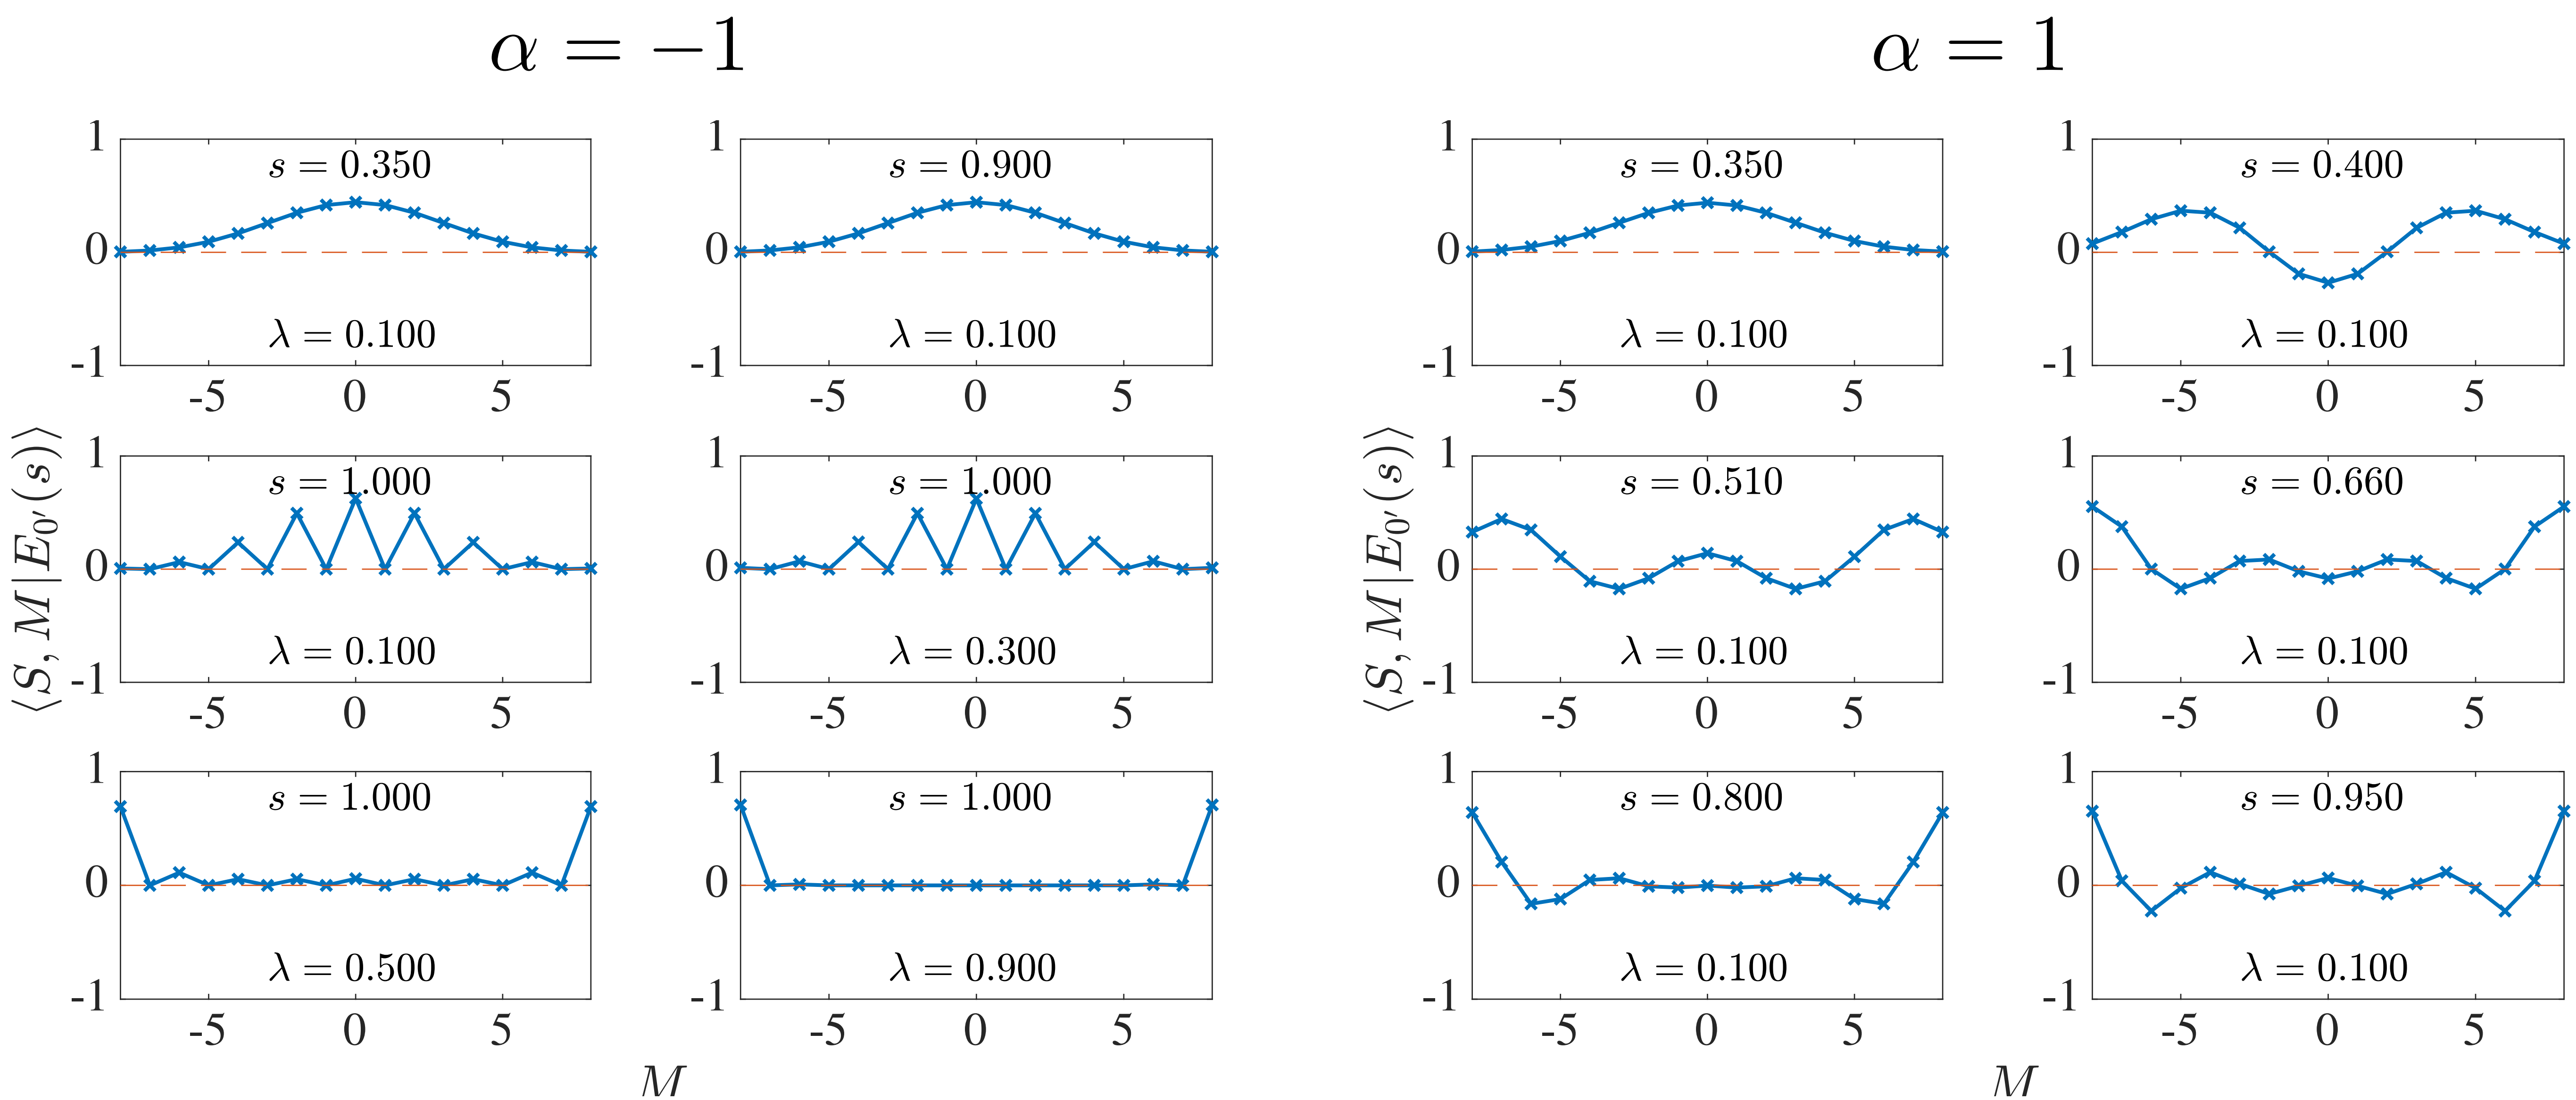

Antiferromagnetic transverse field gives the system more "room to spread"

(1) TA, arXiv:1811.09980.

A Geometrically Local Example

$$H(s) = -(1-s)\sum_i \sigma_i^x + \alpha s(1-s)\sum_{\langle i,j\rangle} \sigma_i^x \sigma_j^x + \frac{s}{6}\sum_{\langle i,j\rangle} J_{ij}\sigma_i^z \sigma_j^z$$

Identical ferromagnetic rings, coupled at their ends

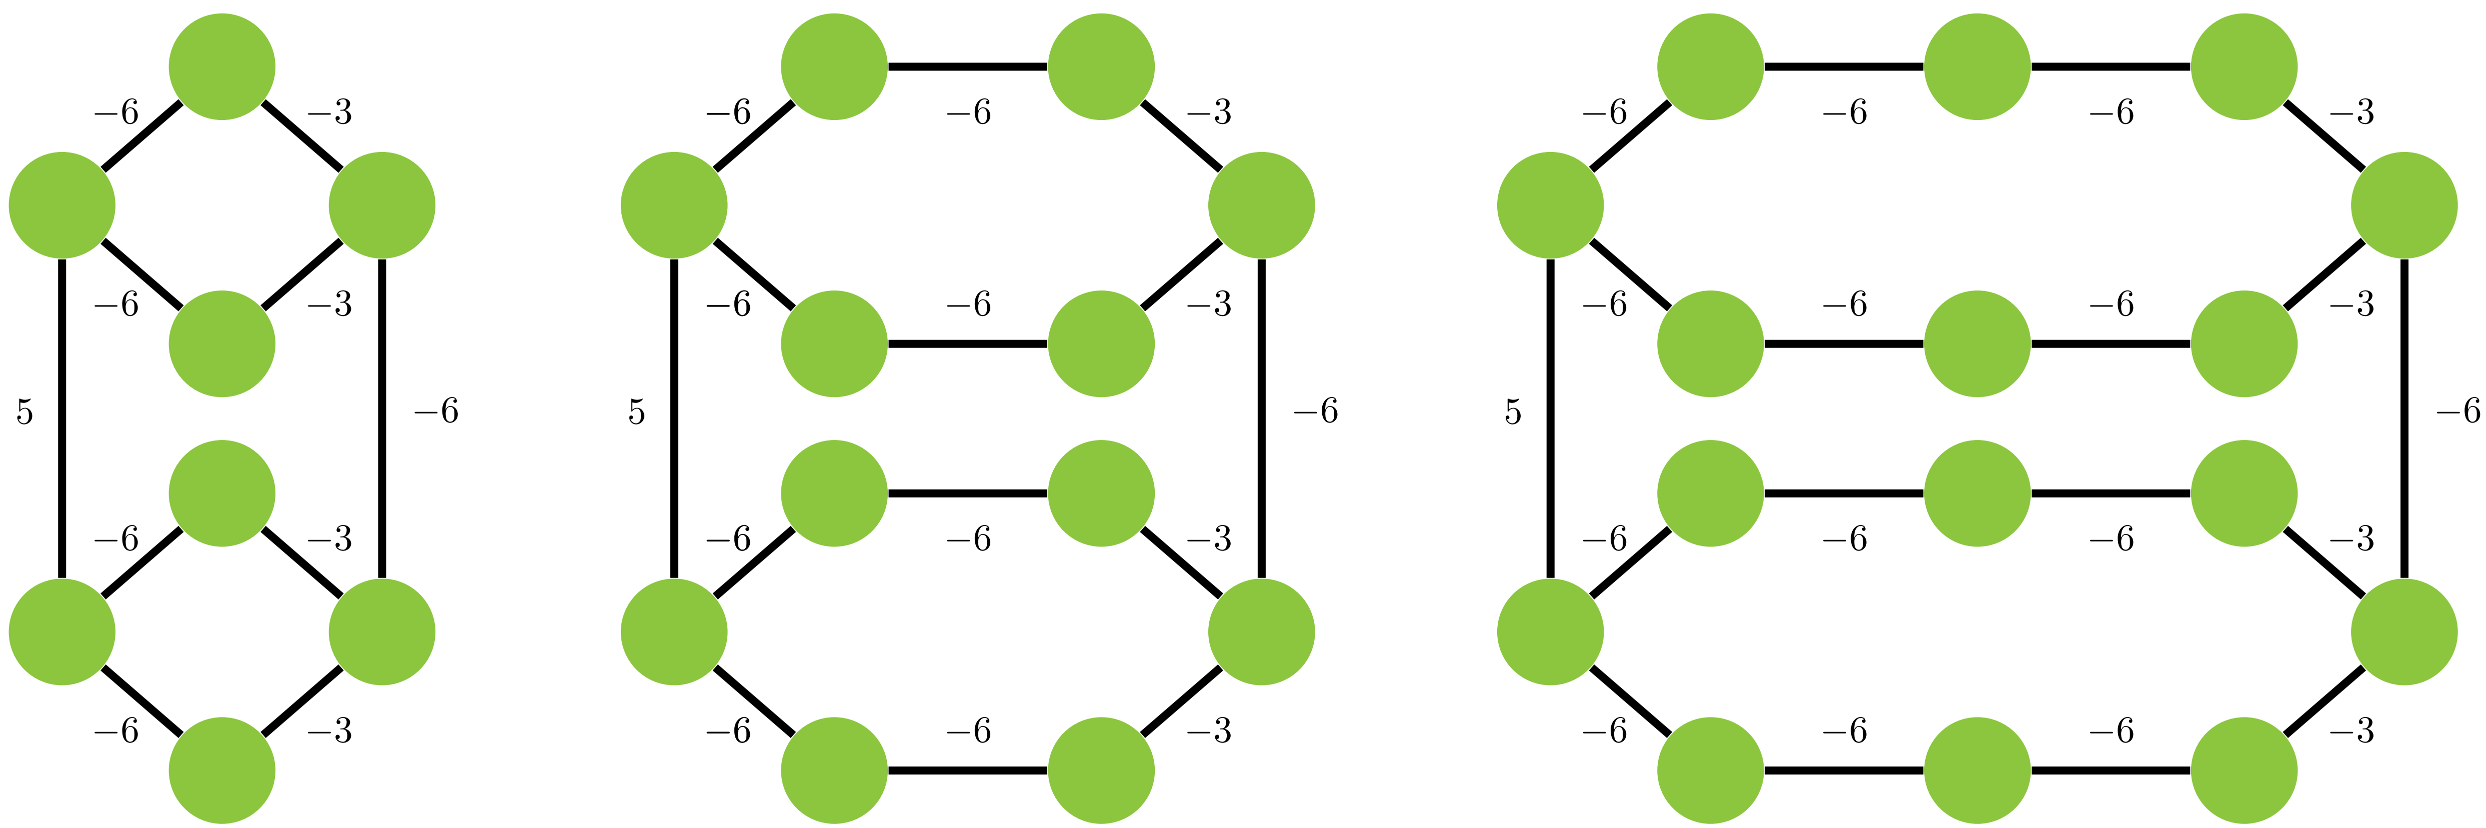

(1) TA, arXiv:1811.09980.

A Geometrically Local Example

$$H(s) = -(1-s)\sum_i \sigma_i^x + \alpha s(1-s)\sum_{\langle i,j\rangle} \sigma_i^x\sigma_j^x + \frac{s}{6}\sum_{\langle i,j\rangle} J_{ij}\sigma_i^z\sigma_j^z$$

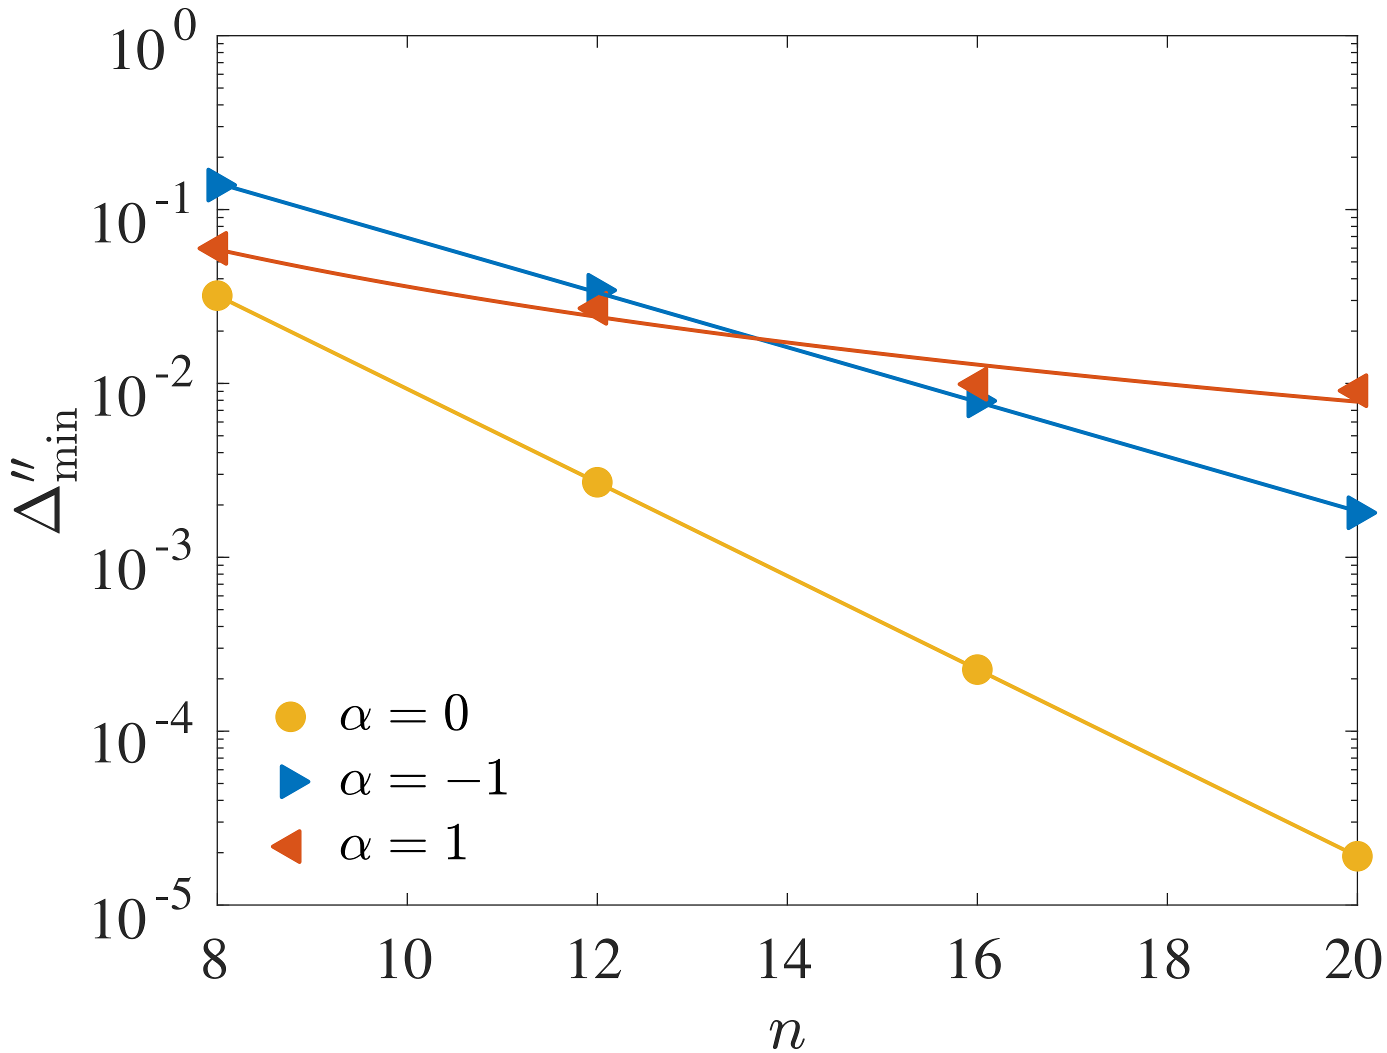

Exponential scaling of the gap for $\alpha = 0, -1$,
but looks like polynomial scaling for $\alpha = 1$

(1) TA, arXiv:1811.09980.

Beyond Standard QA Recap

More exotic interactions open up unexplored parameter spaces for QA

Catalyst Hamiltonians

$$H_{\mathrm{c}} = \alpha \sum_{\langle i,j \rangle} \sigma_i^x \sigma_j^x$$

Beyond 2-body Ising Hamiltonions

$$\sigma_i^z \sigma_j^z \sigma_k^z \qquad \sigma_i^z \sigma_j^z \sigma_k^z \sigma_l^z$$

New annealing protocols like 'reverse annealing' (1-3)

(1) A. Perdomo-Ortiz et al., QIP 10, 33 (2011).
(2) M. Ohkuwa et al., Phys. Rev. A 98, 022314 (2018)
(3) D. Venturelli and A. Kondratyev, arXiv:1810.08584.

Conclusions

Most of our analysis has been restricted to ‘trivial’ computational problems, where analytic/numerical progress can be made.

Will this translate to real world advantages?
Is QA doomed? Is ending on a classical Hamiltonian the problem?
Will we need to adopt universal AQC to see an advantage?

We may be in a situation where the only way to find out is to build such devices and try.

We are living in exciting times;
quantum information processing devices are coming online,
albeit noisy ones, and we have the opportunity to ask:
what can do we do with such quantum computing devices?