

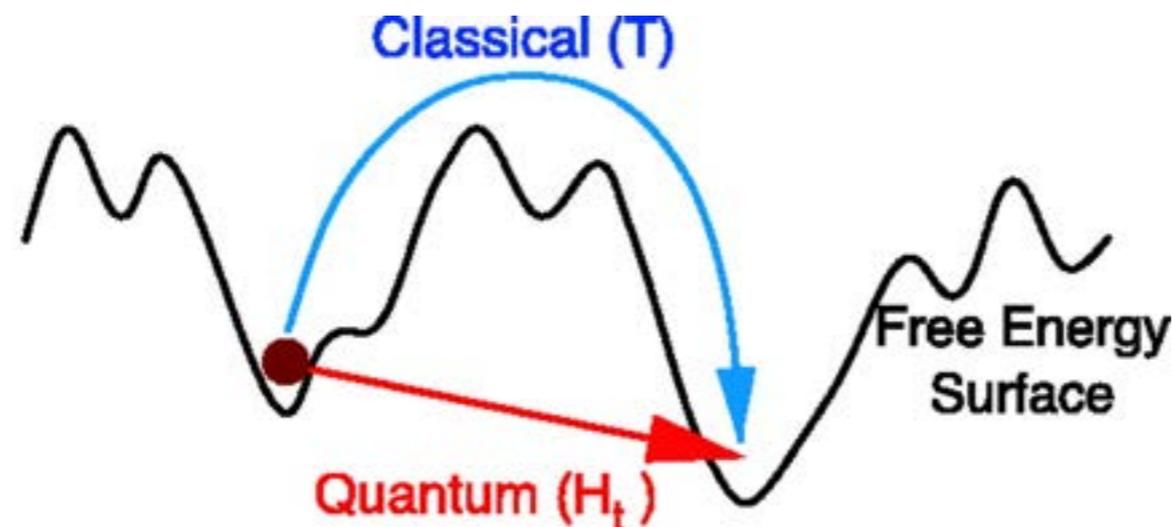
Simulation support for QEO

Matthias Troyer



The big questions to be answered by simulations

- Which problems are amenable to quantum acceleration?
- How can we predict quantum speedup potential?
- Which architecture is best suited to realize quantum speedup?
- What are the requirements on control and calibration?



Picture source: Brooke, Bitko, Rosenbaum, Aeppli, Science (1999)

Simulating quantum annealing

- “Schrödinger” dynamics
 - **Exponential** complexity on classical hardware
 - Simulates the time evolution of a quantum system
 - Unitary evolution in the ground state: **U-QA**
Kadowaki and Nishimori (1998)
 - Open systems dynamics using master equations: **OS-QA**
- Quantum Monte Carlo dynamics (stochastic)
 - **Classical** algorithms with **polynomial** complexity
 - QMC samples the equilibrium thermal state of a quantum system
 - Typically based on path integral Monte Carlo simulations: **QMC-QA**
Apolloni et al (1988), Santoro et al (2002)
 - Mean-field MC version using coherence but no entanglement: **MC-QA**
Shin, Smolin, Smith, Vazirani, arXiv:1401.7087

The case for a software framework

- We are running into a big data problem:
 - $10^3 - 10^4$ instances per problem size
 - $10^3 - 10^9$ repetitions of annealing per problems
 - $O(10)$ different system sizes
 - $O(10)$ different annealing schedules
 - $O(10)$ different annealing times
- We perform more than $O(10^{10})$ simulations
- Job scheduling, launching, and data storage has to be automated to avoid wasting valuable time and resources
- Our framework automates repetitive tasks and substantially decreases the time to obtain results on new problems, algorithms, hardware architectures or annealing strategies

A three-tiered strategy



- We build up our tools and hone our skills starting from small and specifically chosen problems, over medium sized random problems to full applications

Large scale application problems
quantum enhancement in applications

Medium sized engineered problems
explore which random problems show quantum enhancement

Small tailored problems
learn the mechanisms of quantum enhancement

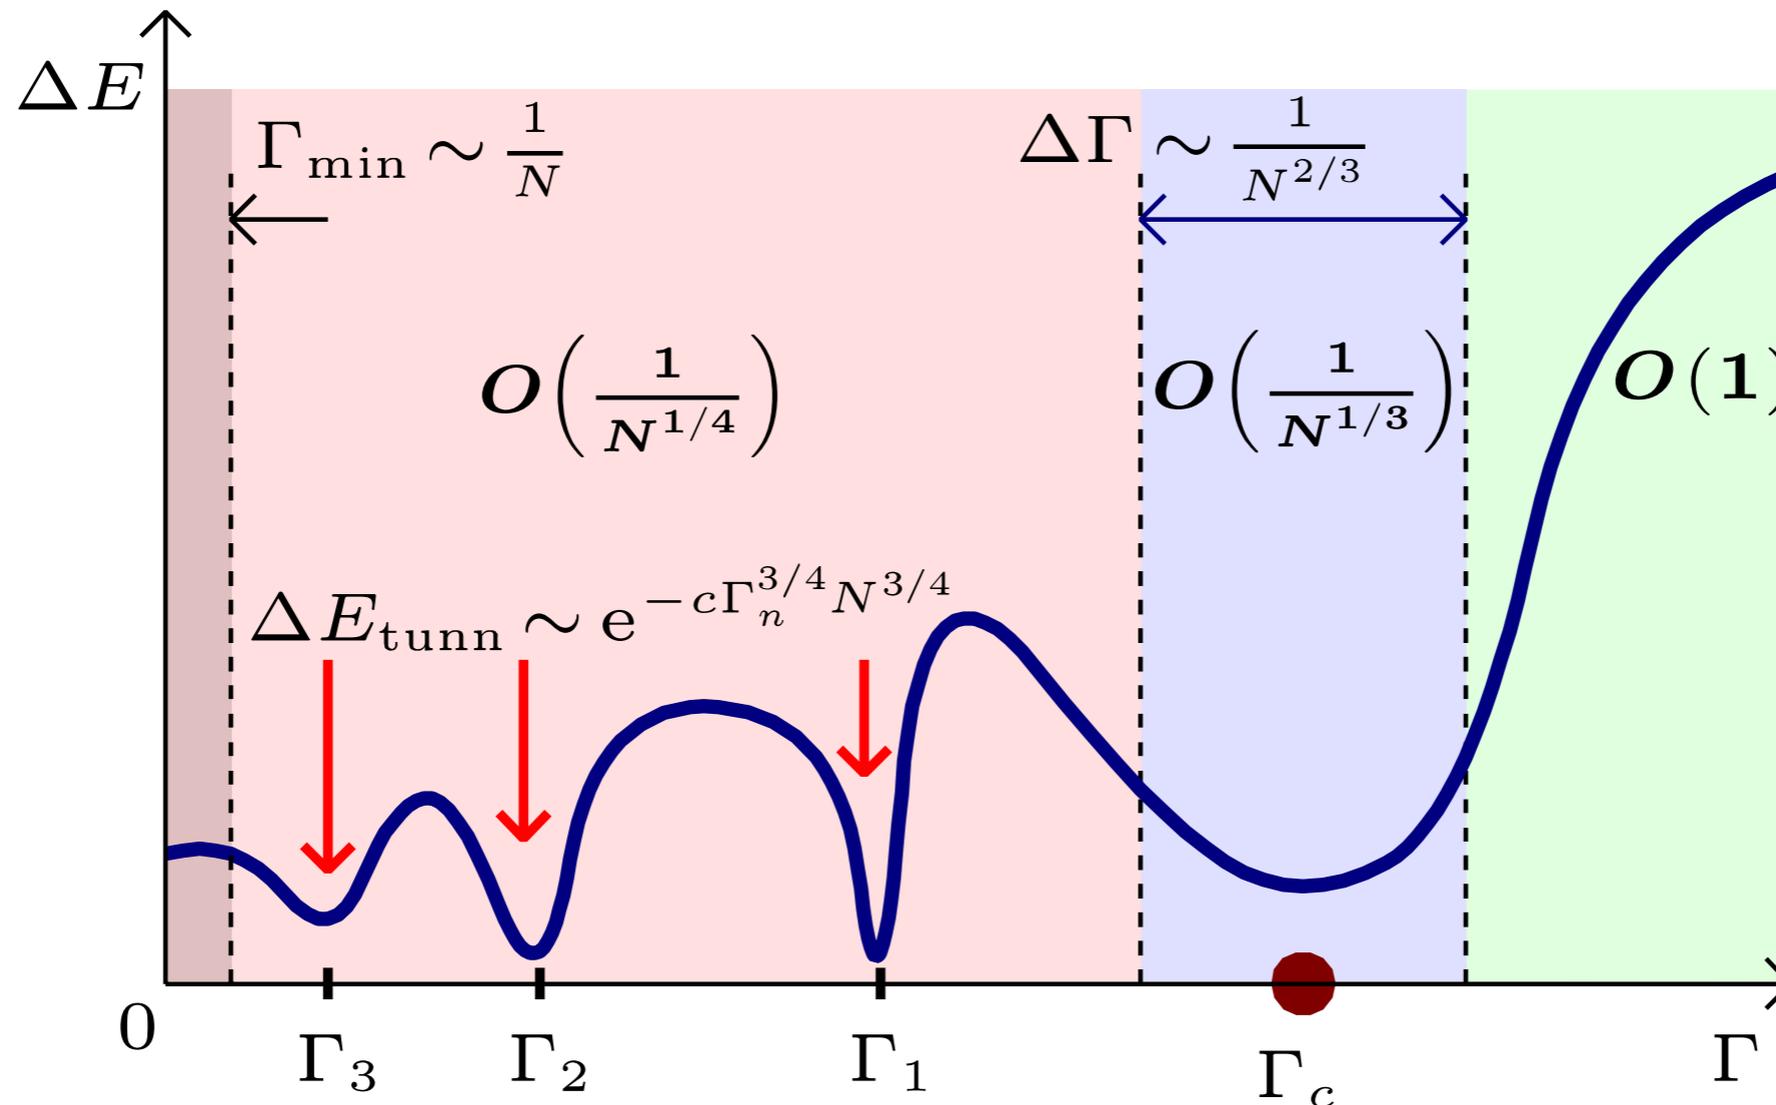
The physics of quantum annealing

Hardness Metric: Density of QA bottlenecks

S. Knysh, arXiv:1506.08608

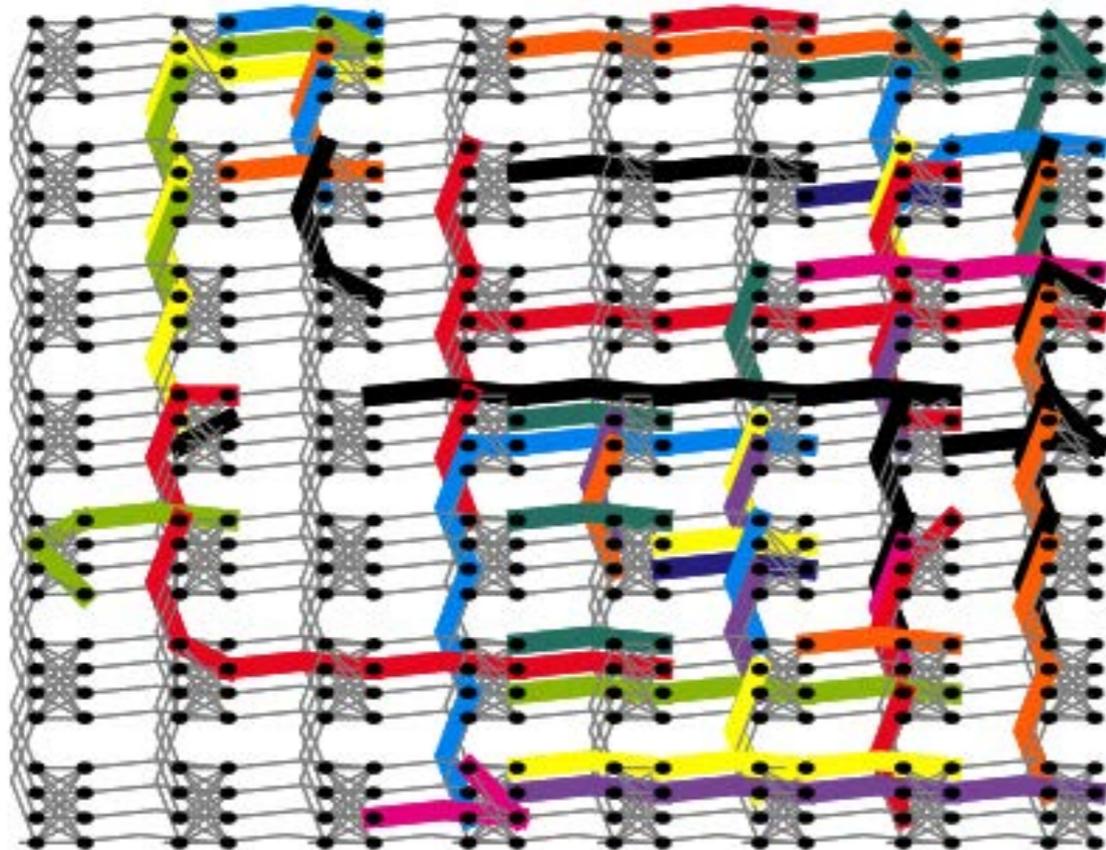
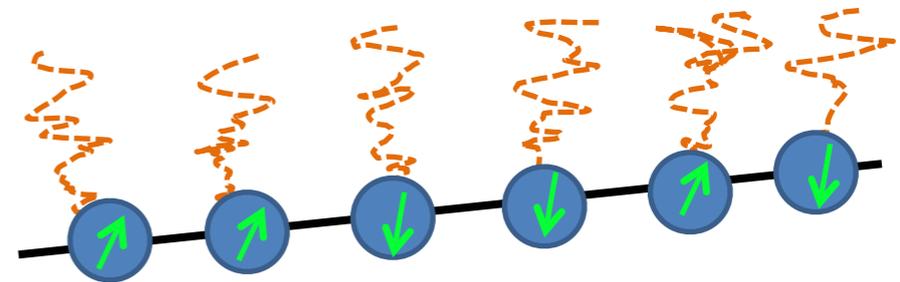
A growing number of exponentially small gaps appears inside the spin glass phase and dominate the hardness

Power law gap at spin glass phase transition



Theory of dissipative QA on linear chains

Linear chains or trees are fundamental primitives and basis for all problem embedding

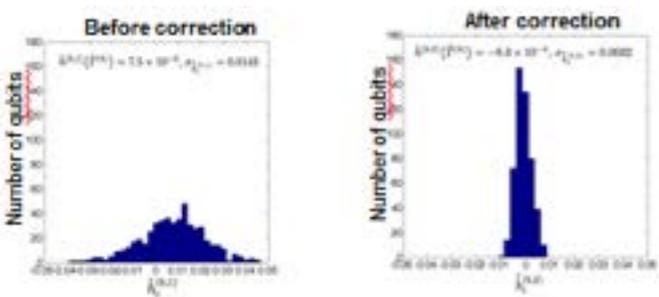


Speedup of open system QA compared to fully coherent annealing

$$v_{\text{opt}}/v_{\text{KZ}} \propto (\beta/\alpha) \ln(\beta/\alpha)^{3/2} \gg 1$$

V.Smelyanskiy, D. Venturelli et al. (2015)

More NASA Ames projects



❑ Baseline QA performance established for **scheduling** and **fault diagnostics** problems

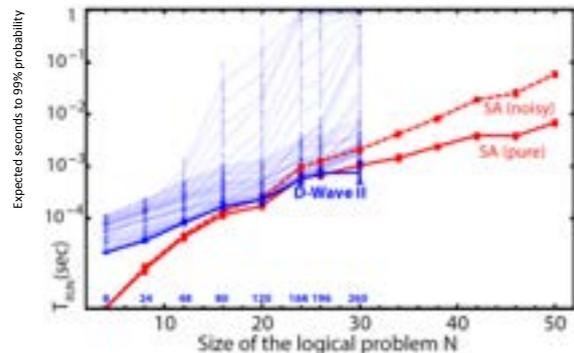
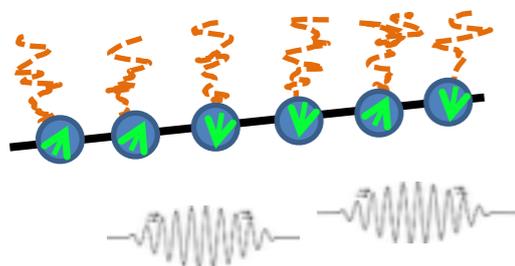
❑ Best strategies for **mapping and embedding**

❑ **Hybrid quantum-classical** optimizers

❑ Develop **device tuning and calibration** techniques to reduce the effect of **misspecification** errors.

❑ Design problems to **unveil and evaluate quantum speedup**

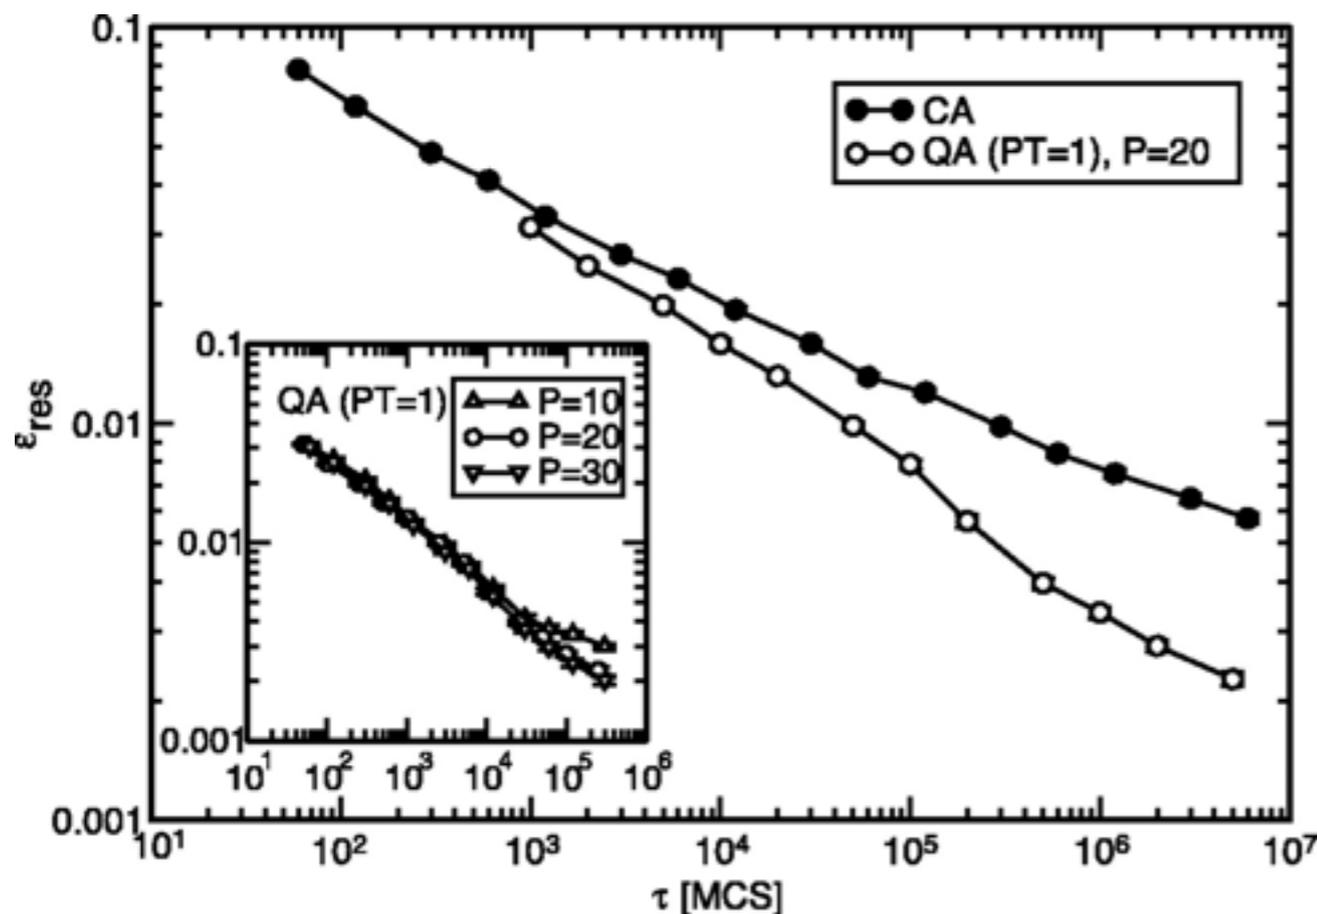
❑ Identify and overcome the **bottlenecks of quantum annealing**



Quantum Monte Carlo annealing

Simulated classical versus quantum annealing

- Early evidence for superiority of QA came from QMC simulations
- Lower residual energies in simulated quantum annealing compared to simulated thermal annealing, but not for all models



Santoro *et al*, Science (2002)

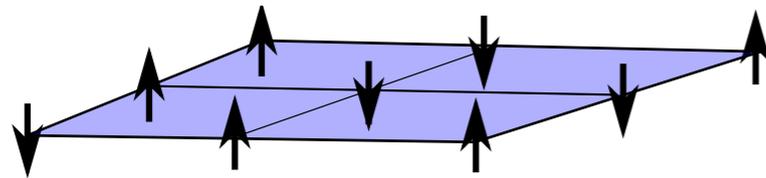
similar results by
Matsuda, Nishimori, Katzgraber (2009)

this is inconsistent with what we saw on D-Wave for 2D spin glasses!

From classical to quantum Monte Carlo

- Simulated annealing performs a Monte Carlo simulation
- Sample random configurations according to their Boltzmann weights

$$Z = \sum_{s_1, \dots, s_N} \exp(-\beta \sum_{i < j} J_{ij} s_i s_j)$$



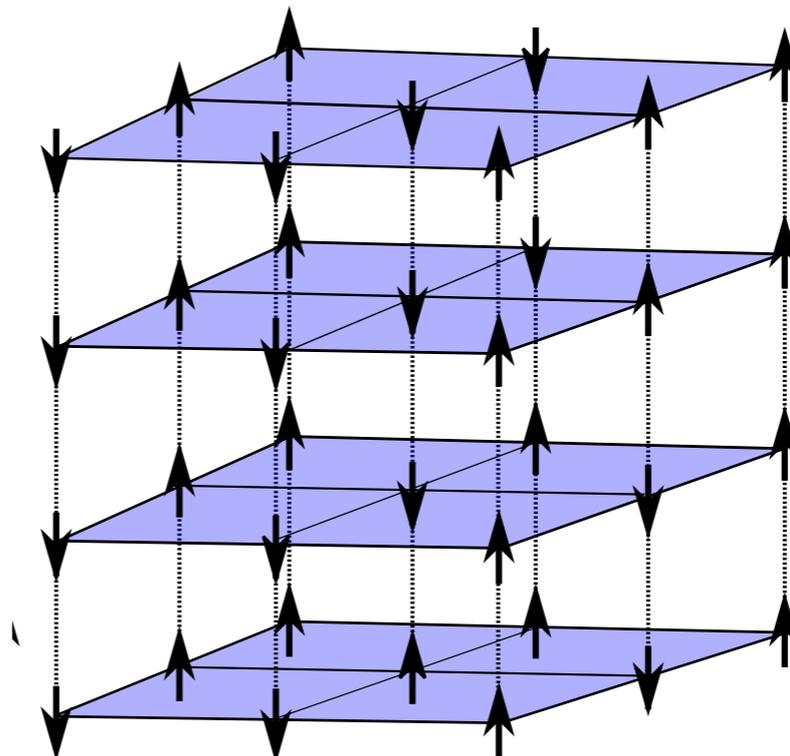
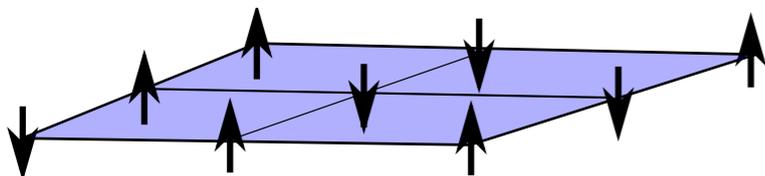
From classical to quantum Monte Carlo

- SQA-QMC performs a **quantum** Monte Carlo simulation
- Map quantum system to classical path integral
- Sample random configurations according to their weights

$$Z = \text{Tr} \exp(-\beta H) \neq \sum_c \exp(-\beta E_c)$$

$$Z = \text{Tr} \exp(-\beta H) = \text{Tr} \exp[(-\Delta_\tau H)^M]$$

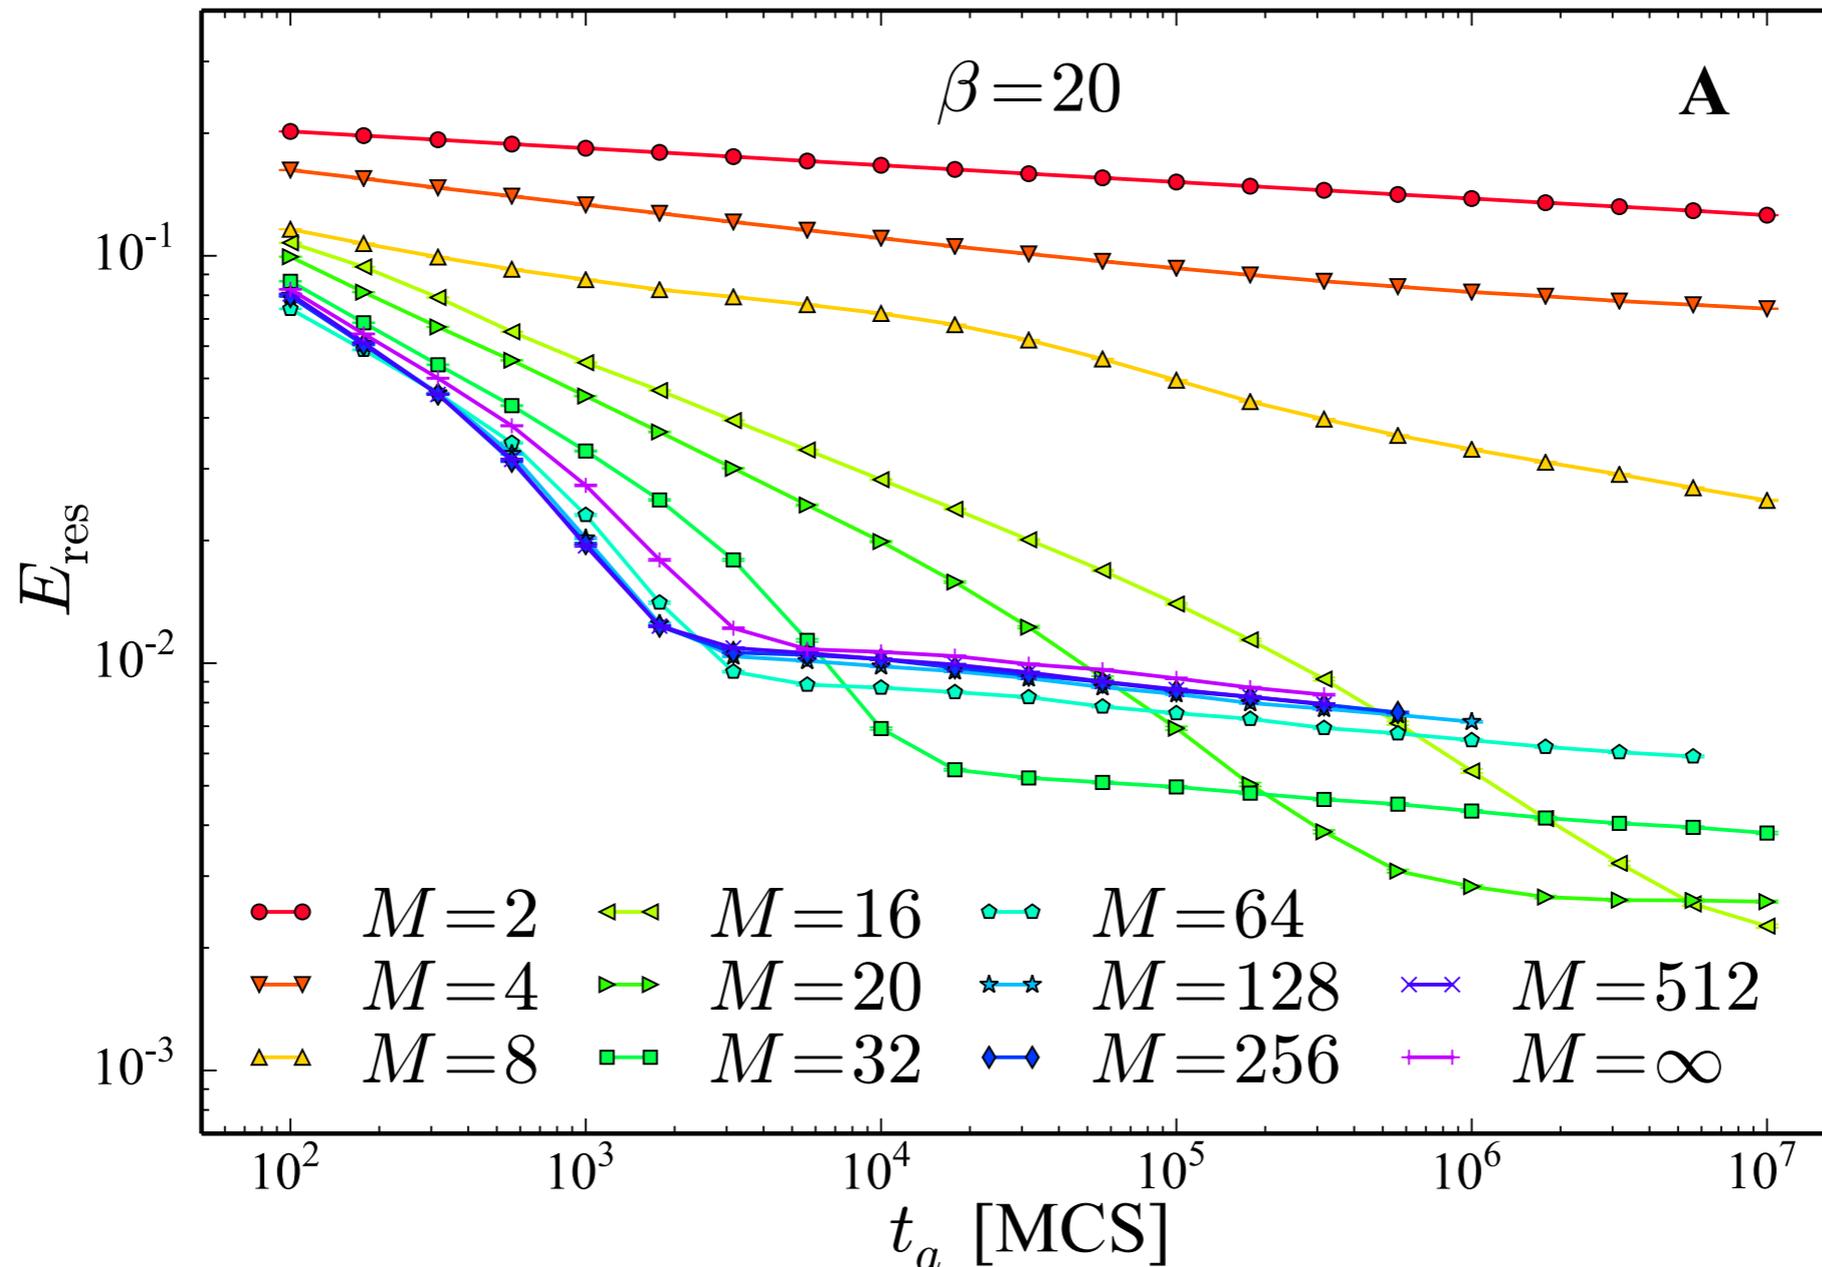
$$\approx \sum_{(s_{i,\tau})} \exp \left(-\Delta_\tau \sum_{i < j} \sum_{\tau} J_{ij} s_{i,\tau} s_{j,\tau} - \Delta_\tau \sum_i \sum_{\tau} J' s_{i,\tau} s_{i,\tau+1} \right) + O(\Delta_\tau^2)$$



Time step dependence

Heim, Rønnow, Isakov, MT, Science (2015)

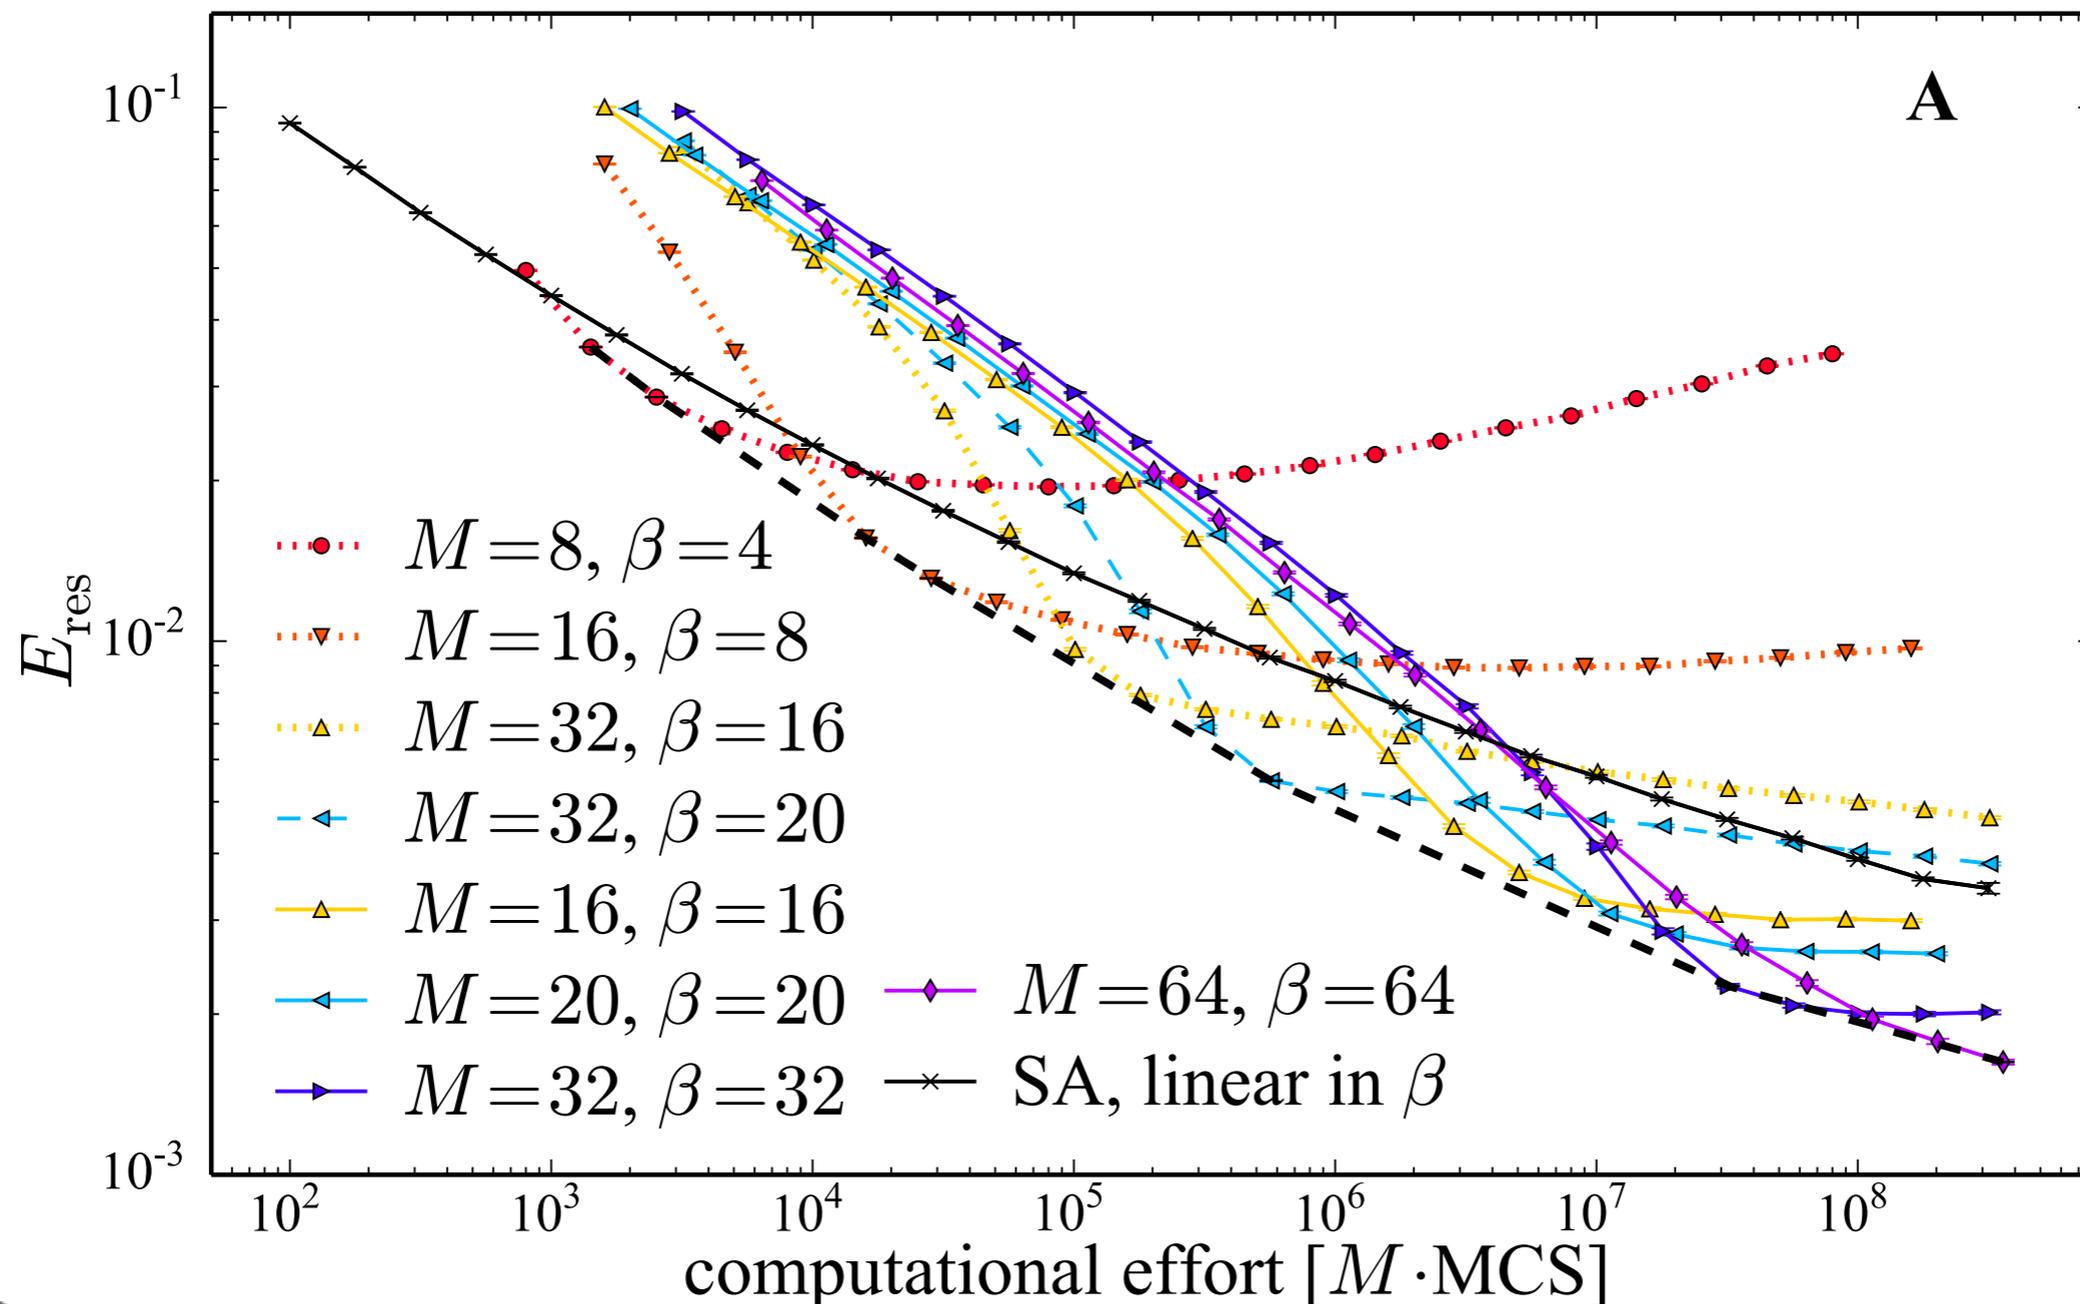
- The behavior of a simulated quantum annealer depends strongly on discretization of imaginary time path integrals



Performance as a classical optimizer

Heim, Rønnow, Isakov, MT, Science (2015)

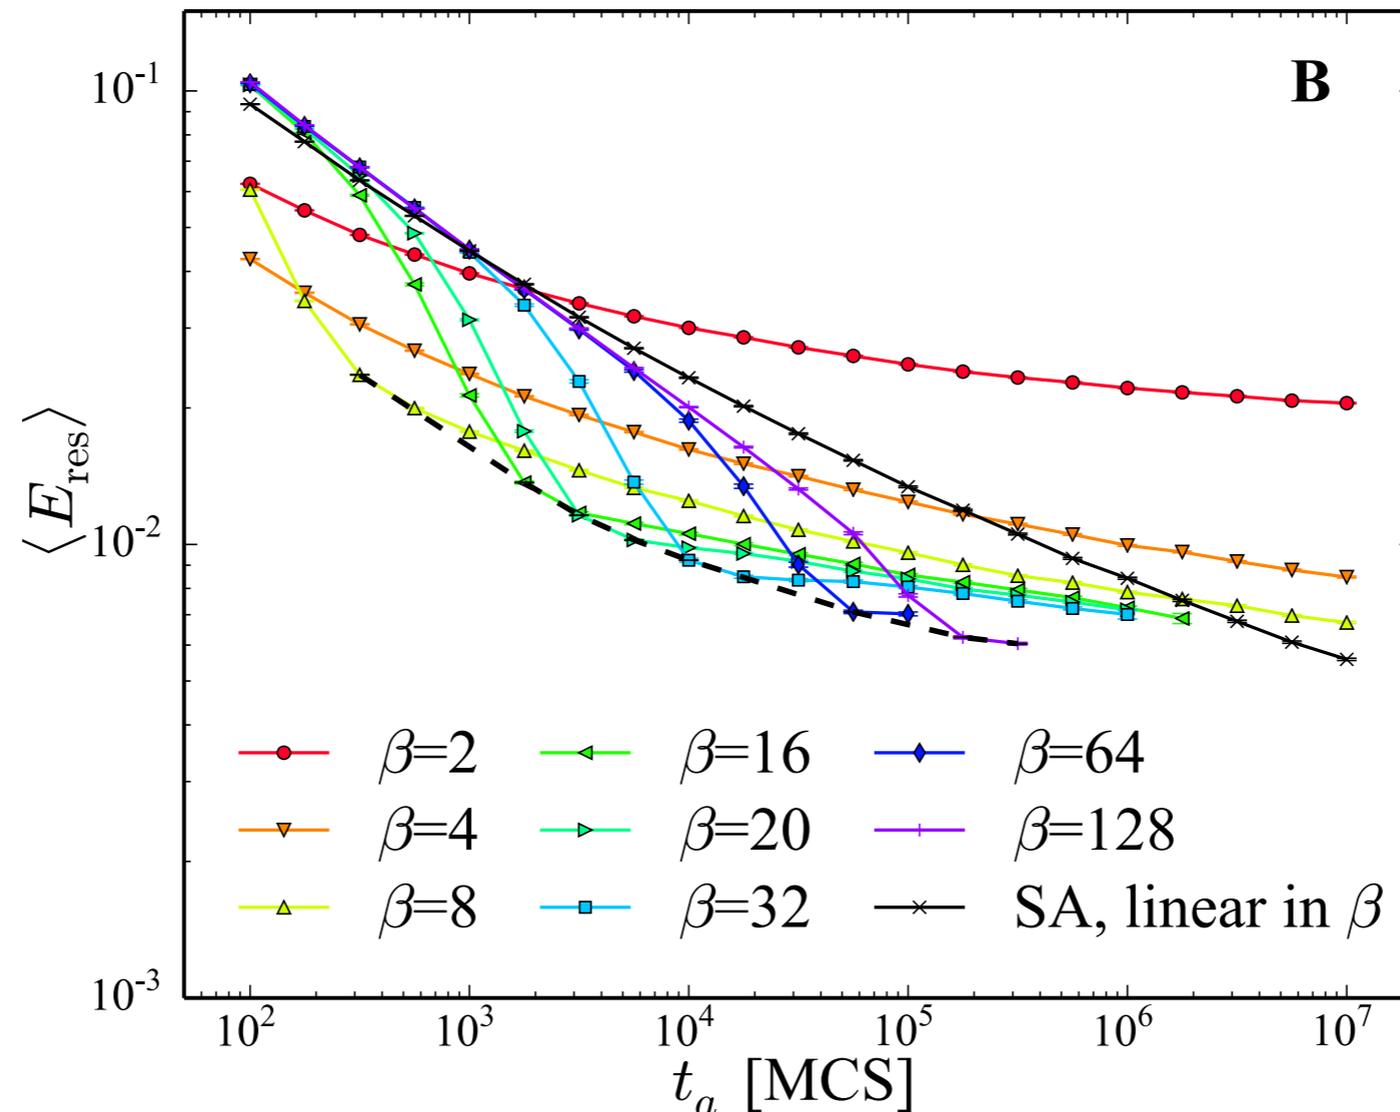
- Use a large time step to get best performance



Performance in the “physical” limit of zero time step

Heim, Rønnow, Isakov, MT, Science (2015)

- One should use infinitesimal time steps to make predictions for devices
- Quantum annealing quickly finds low energy solution
- But in this case ultimately scales worse than simulated annealing



Lessons learned for quantum annealing

Heim, Rønnow, Isakov, MT, Science (2015)

- QMC annealing can be performed in different ways
- “Cheat” with large time steps in the path integral to use QMC as a quantum-inspired classical algorithm
- Work in continuous time to mimic a physical quantum annealer and assess potential for quantum speedup

Are quasi-2D spin glasses a bad choice?

- 2D spin glasses

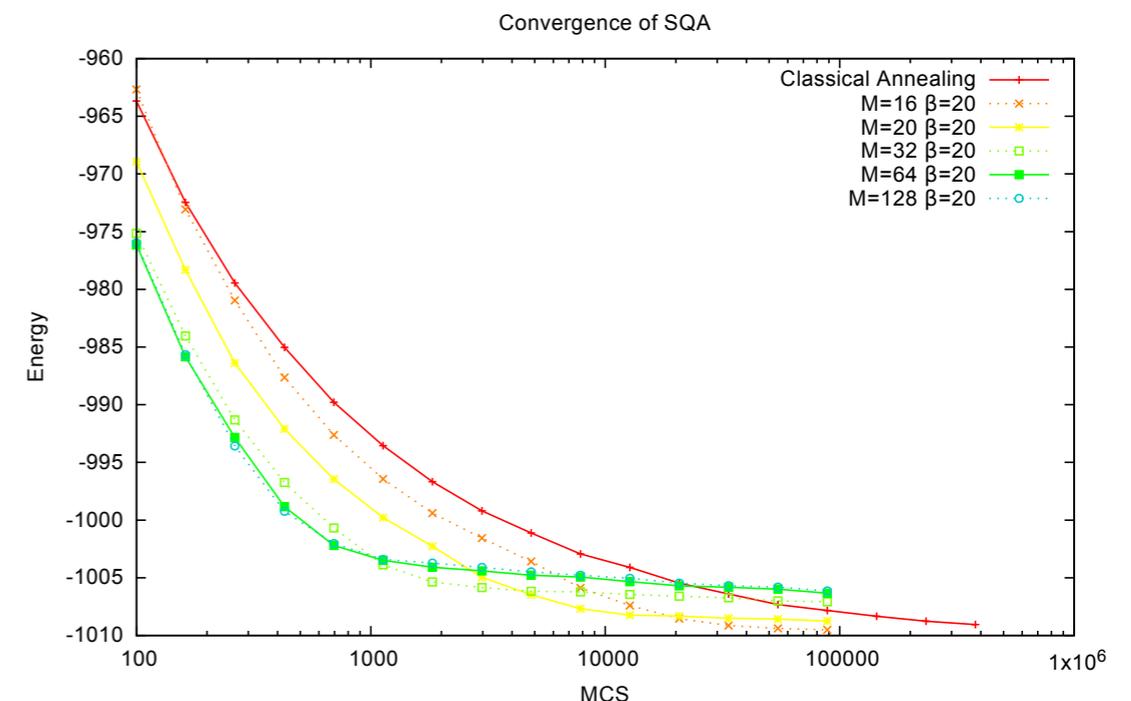
- Critical temperature $T_c=0$
- Infinitesimally small temperature lets the system escape a local minimum
- The barriers may be shallow and not amenable to quantum speedup

Katzgraber *et al*, PRX (2014)

- 3D spin glasses

- Critical temperature $T_c \approx 1$
- Barriers are extensively high
- But similar behavior

- We need to engineer harder problems!



Engineered random instances

Create tunable, sensitive probes for quantum enhancement.

Engineering harder instances

arXiv:1505.01545

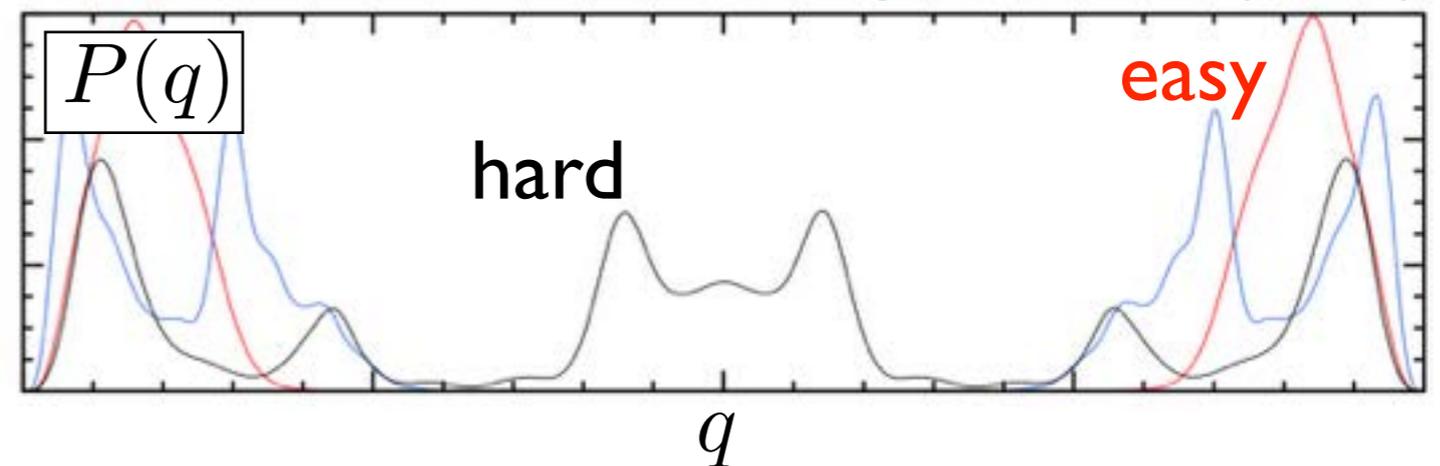
- Reduce the degeneracy with Sidon sets. Example:

$$Q_{ij} \in \{\pm 1\} \longrightarrow Q_{ij} \in \{\pm 5, \pm 6, \pm 7\}$$

- Find tunable instances:

- Study the distribution of the order parameter $P(q)$.
- Many peaks means many metastable states (hard).

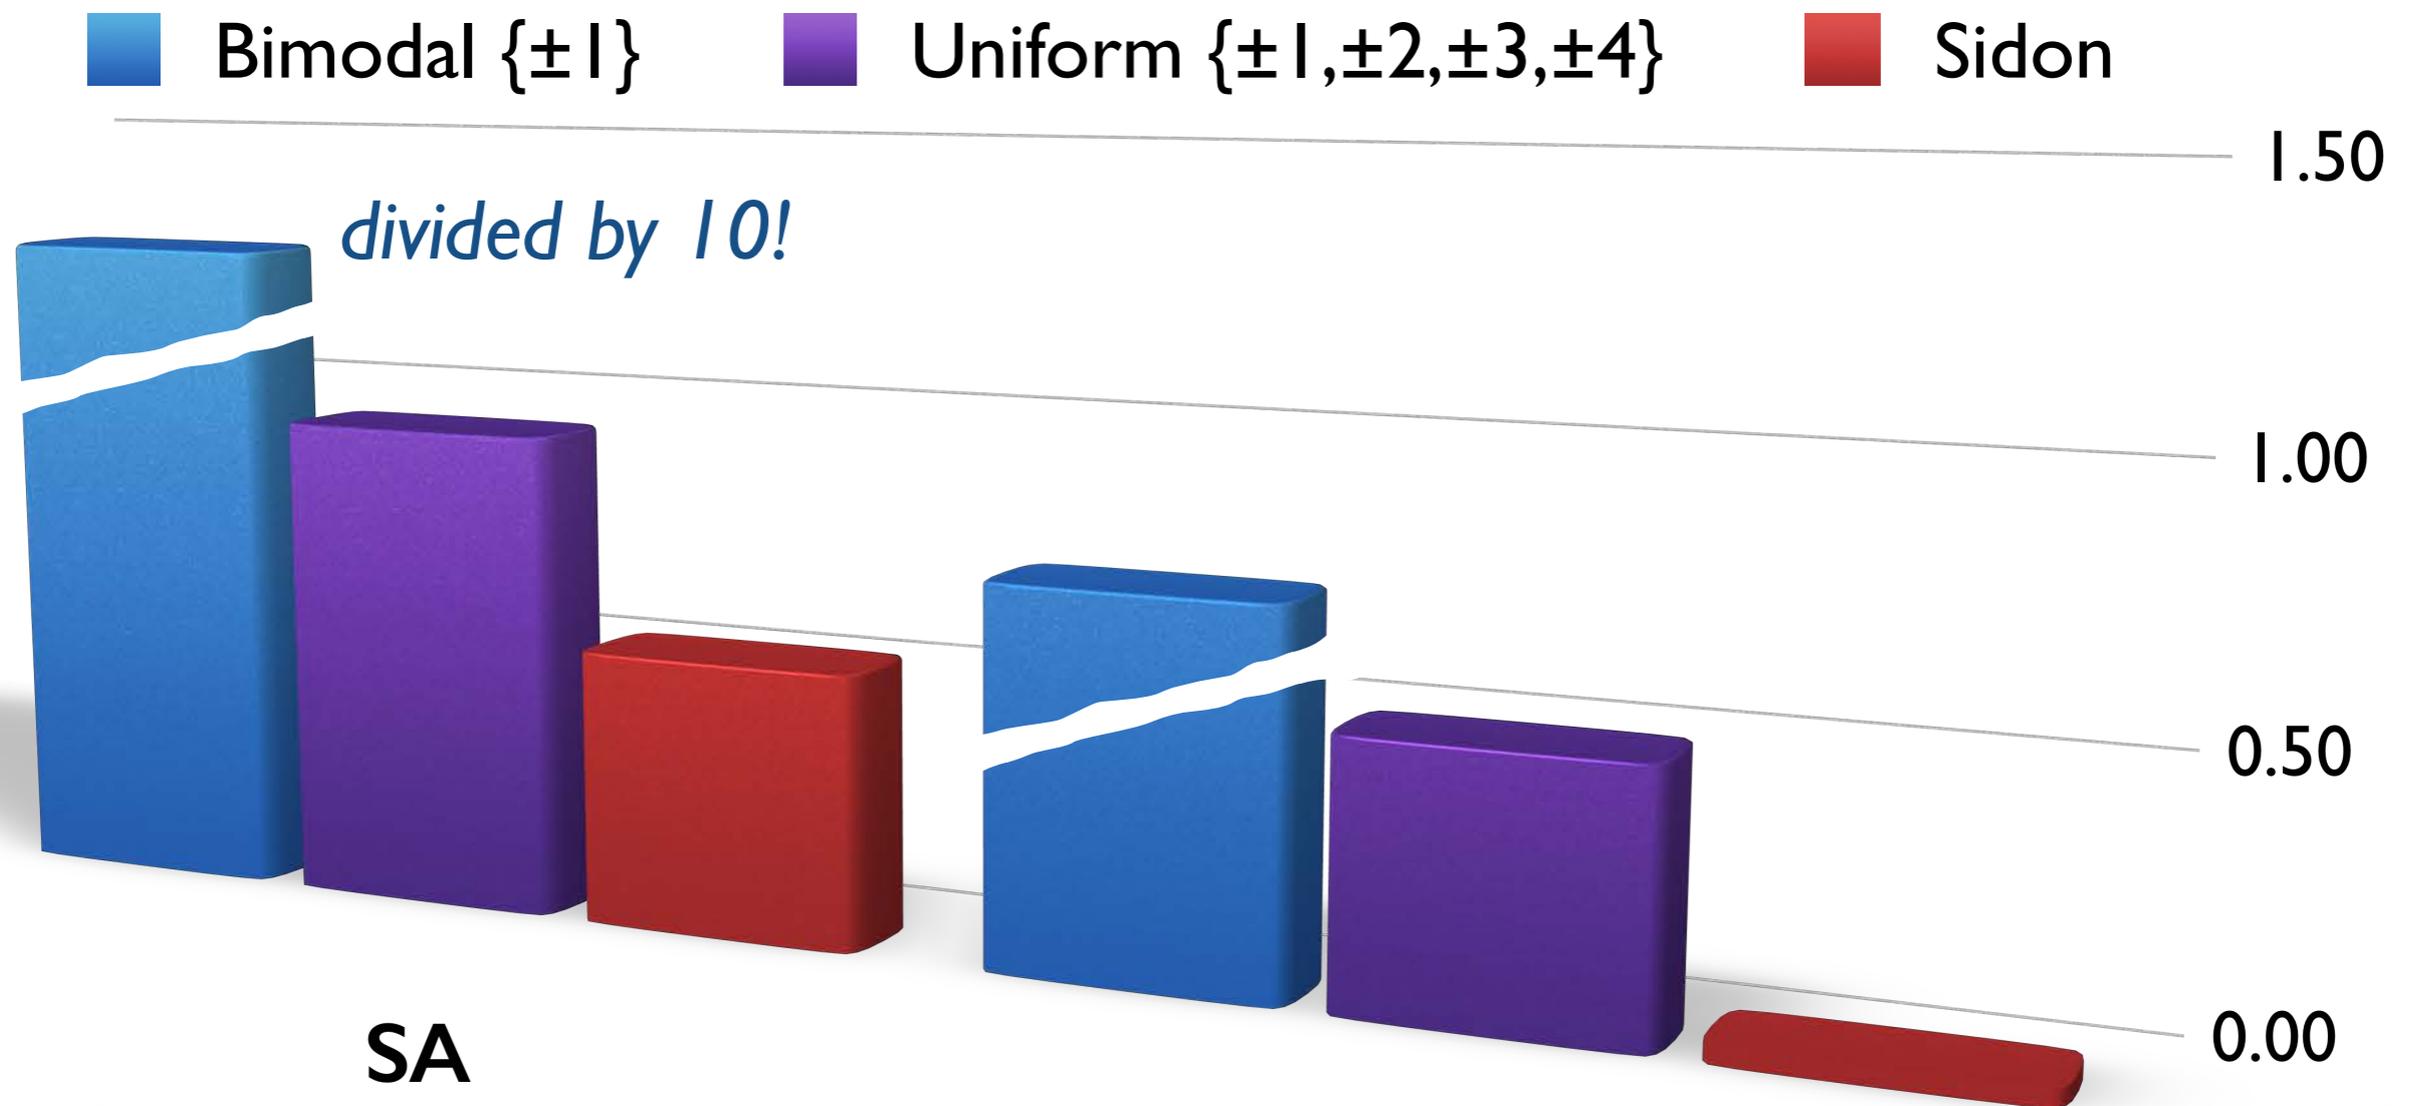
Yucesoy et al. PRL (2012)



- More general: Post- & pre-processing:

- Multi-tier physics-based data filtering.
- Novel planting approaches.
- Extensive data mining.

Example *average* success probability in %



SA

DW2

easier

harder

*Statistical physics
is key in the design!*

Hardening random instances

- **Algorithm to increase number of unique GS instances:**

- Assign bond couplings randomly
- If a spin has Z even neighbors, choose the last bond coupling such that the local field is never zero *for all* possible spin configurations.

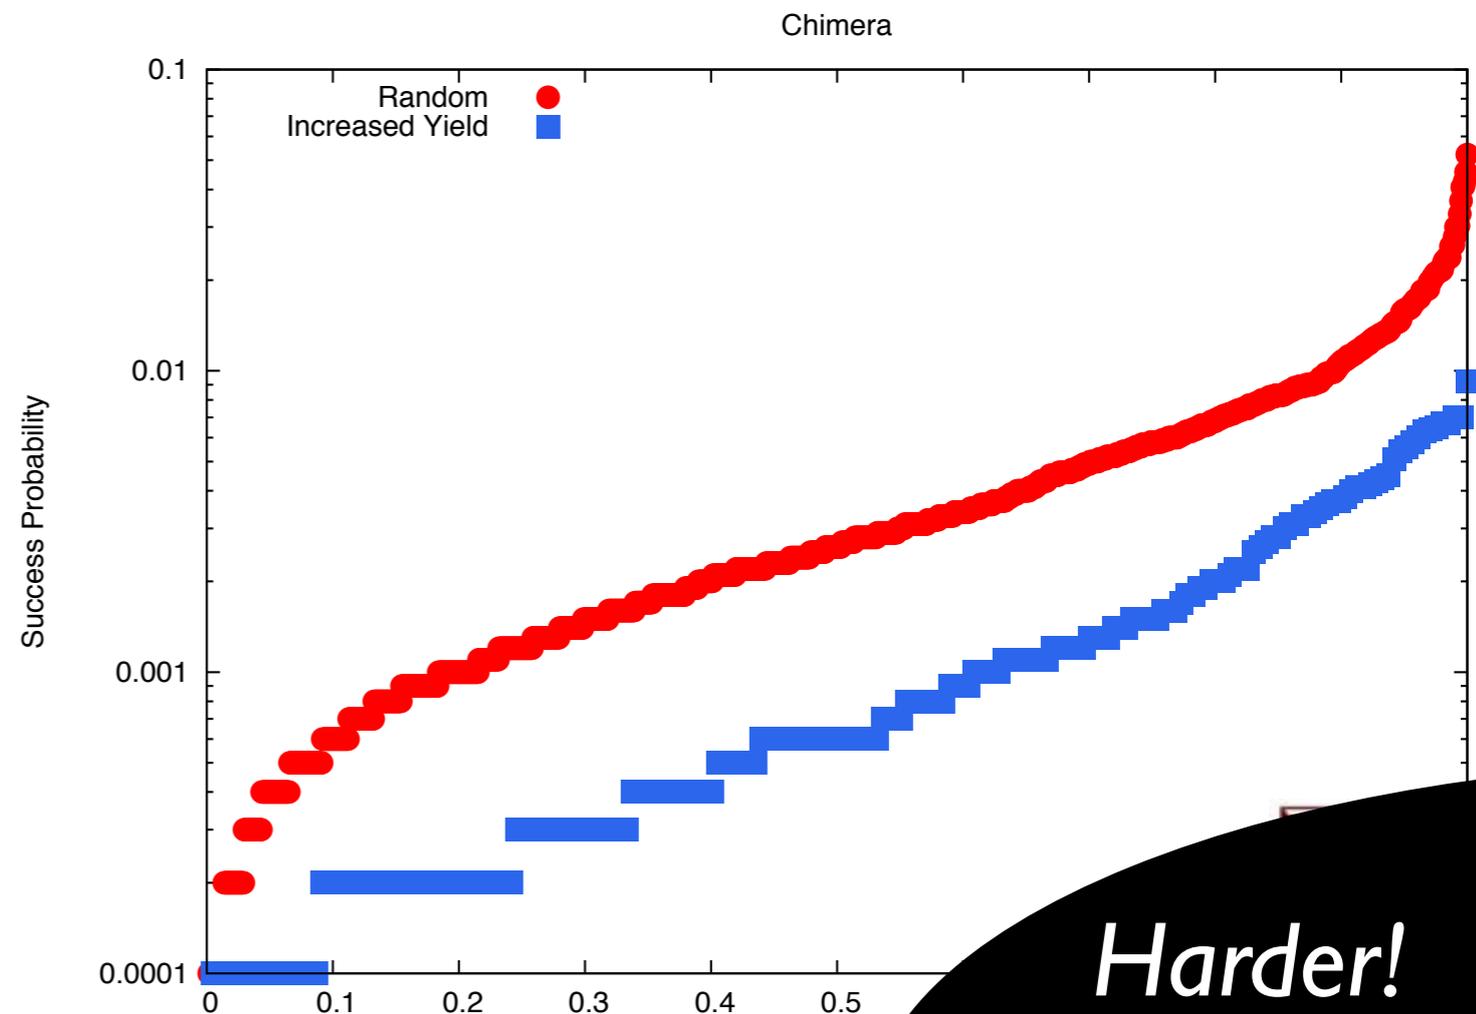
- **Results:**

- 4 times more unique GS.

$$\mathcal{Y}_{\text{random}} = 4.5\%$$



$$\mathcal{Y}_{\text{new}} = 20.8\%$$

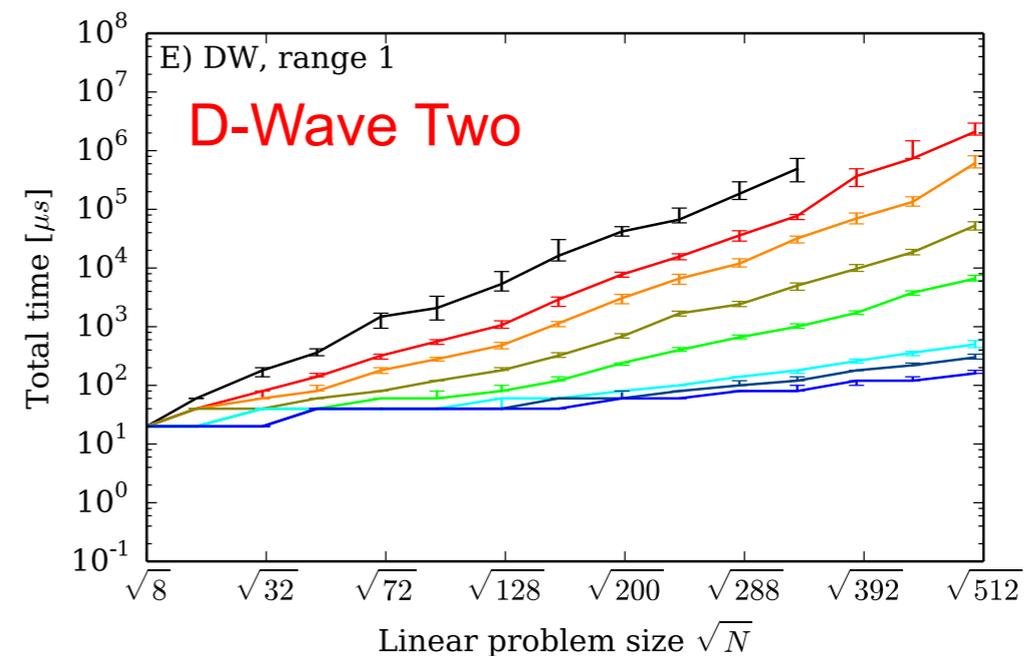
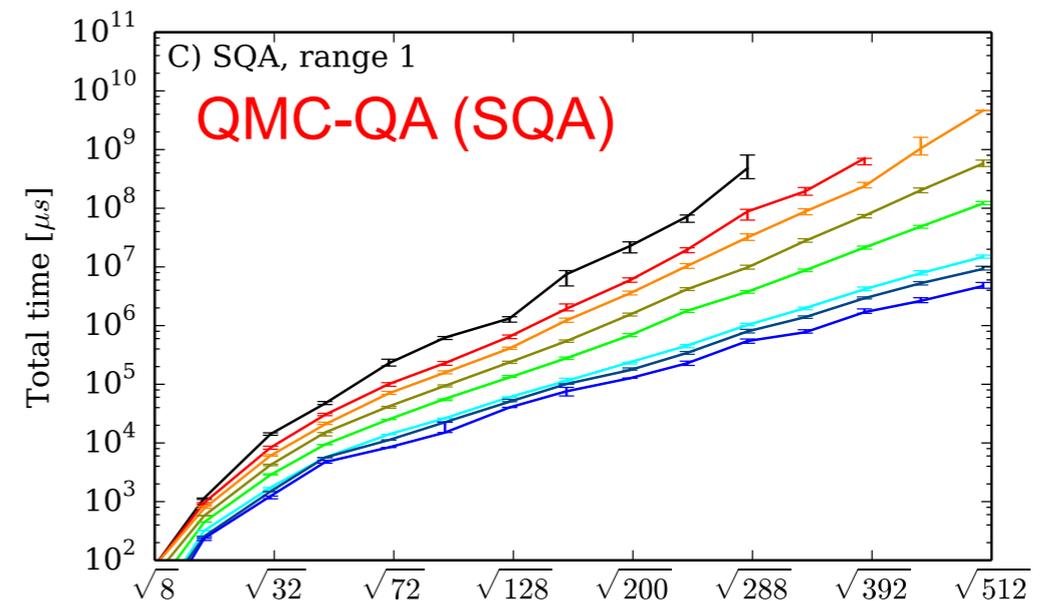
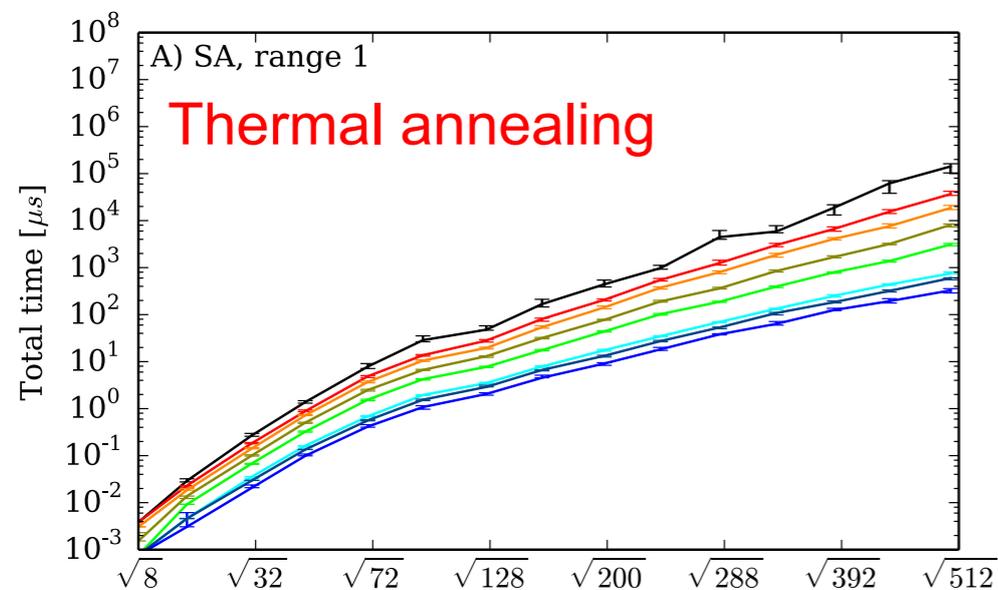


Harder!

Heavy tails ...

Look at more than the median (typical) scaling

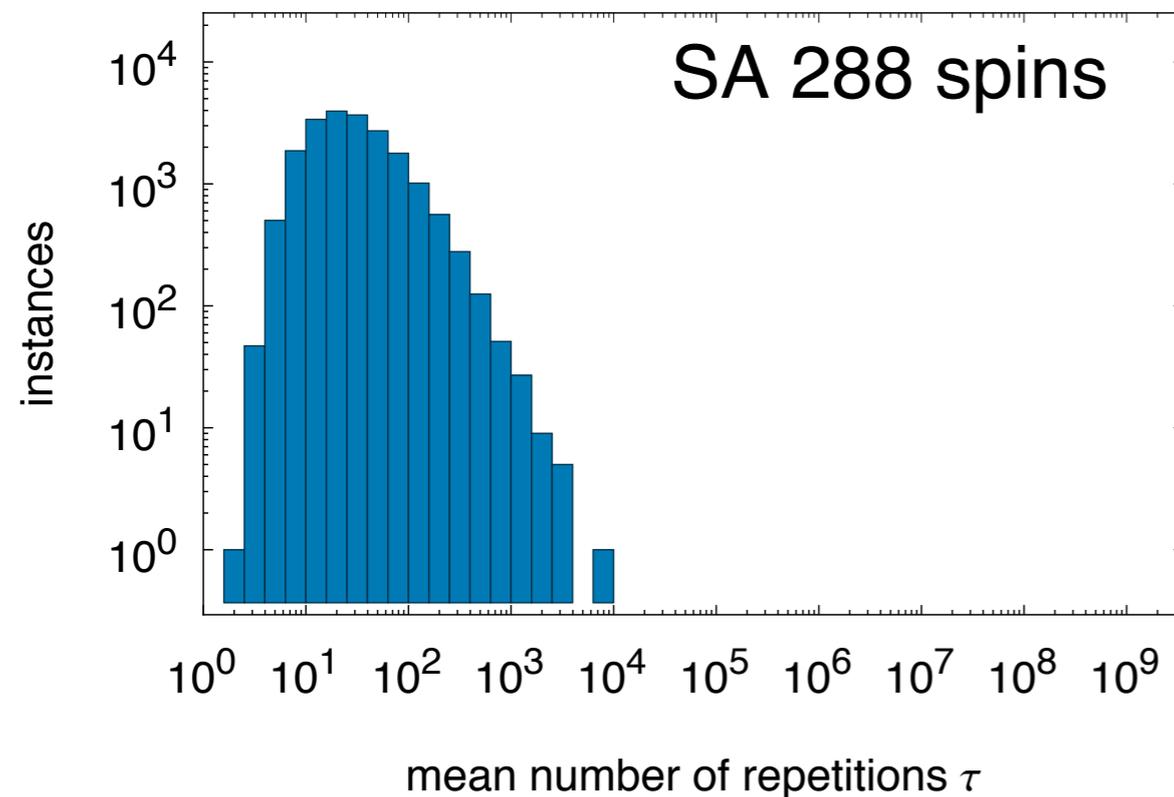
- Time to solution varies hugely between instances
- Quantum annealing seems to have more problems for high quantiles (hard instances)



Heavy tails

- Distribution of time to find a ground state has fat (power law) tails
- Tails are fatter for quantum annealing than for thermal annealing

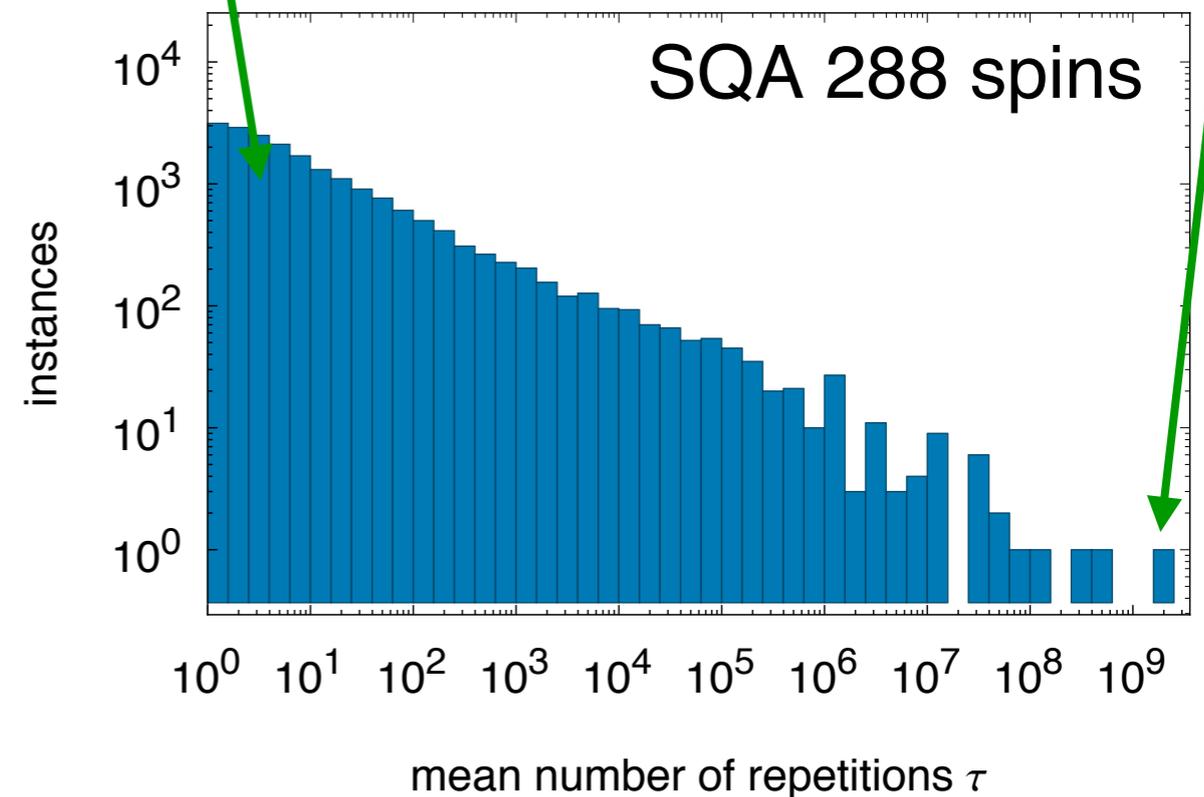
simulated classical annealing



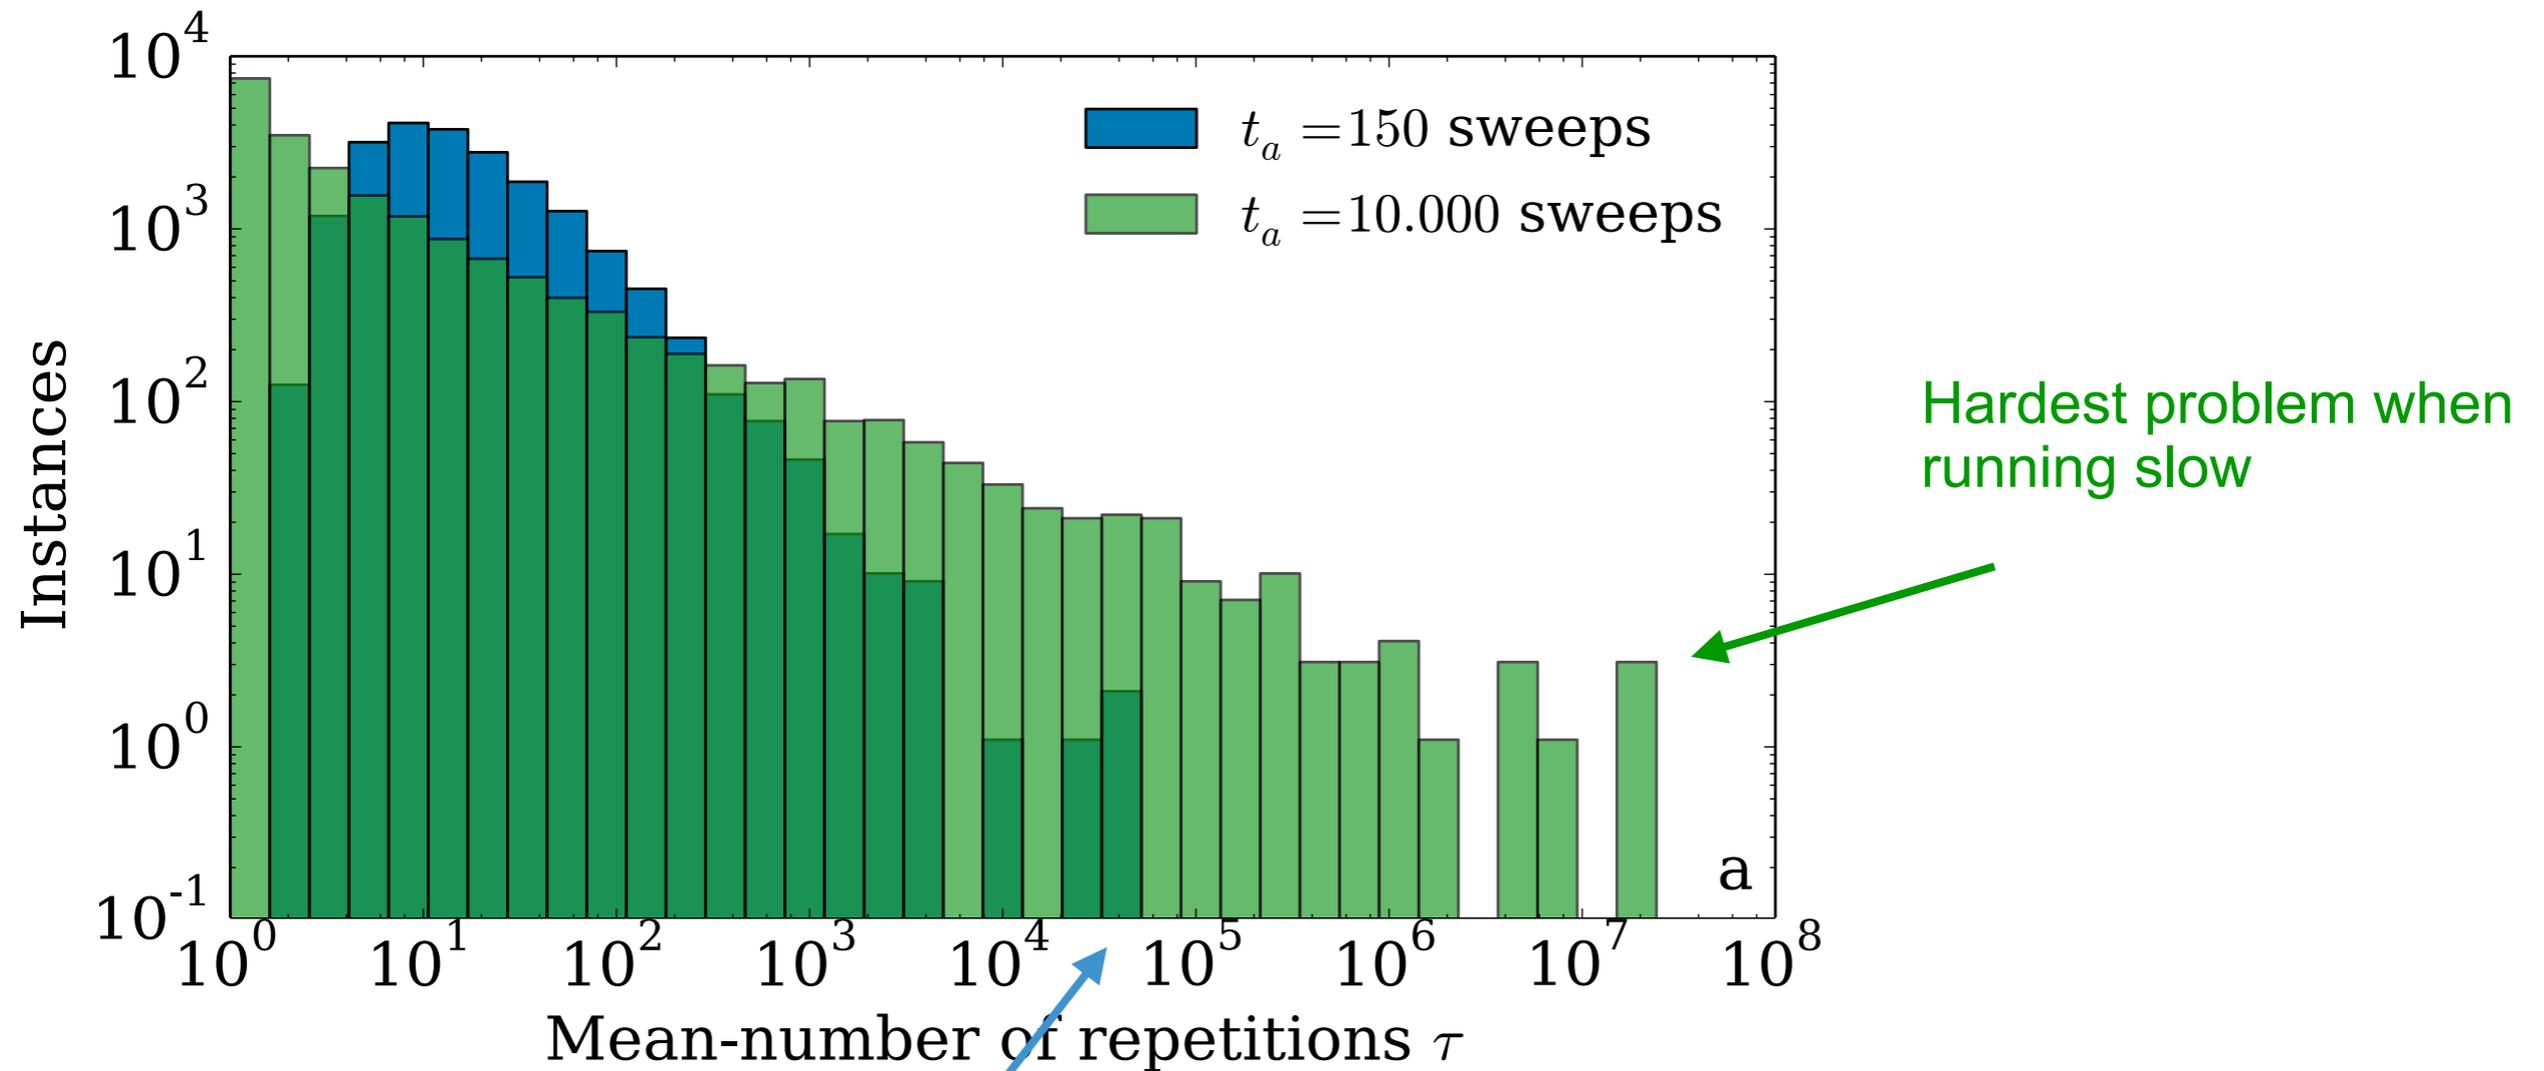
Easy problems

Hardest problem

simulated quantum annealing



Annealing faster is better for hard problems



**Speedup of 45'866 when running hardest problems faster:
66.6 times faster needs 688 times less repetitions**

Watch the tails!

- A large fraction of “rare events” takes orders of magnitude longer than typical instances
- Quantum annealing on chimera instances shows **much** fatter tails than simulated annealing.
- Optimizing the annealing schedule, and especially the annealing time gives improvements by many orders of magnitude!
- Slow (adiabatic) annealing can be sub-optimal. Much better performance on hard problem instances is seen when annealing faster

But does QMC tell us anything about the real performance of a quantum annealer?

Mazzola, Troyer (ETHZ)

Isakov, Smelyanskiy, Boixo, Neven (Google)

Zhang Jiang (NASA)

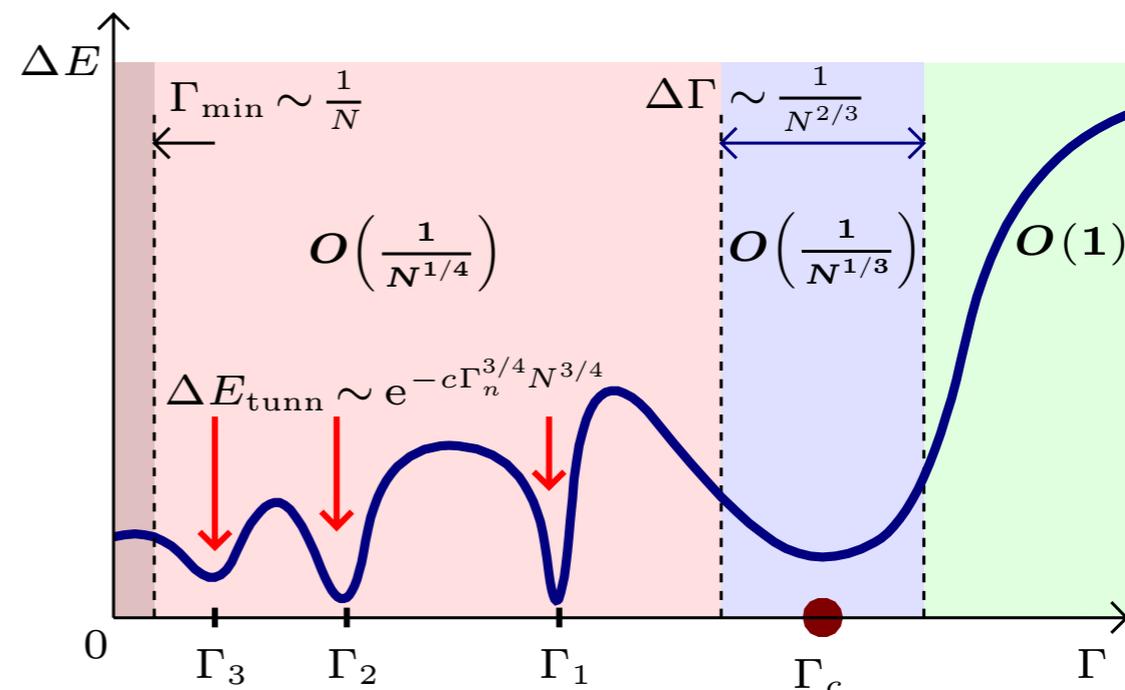
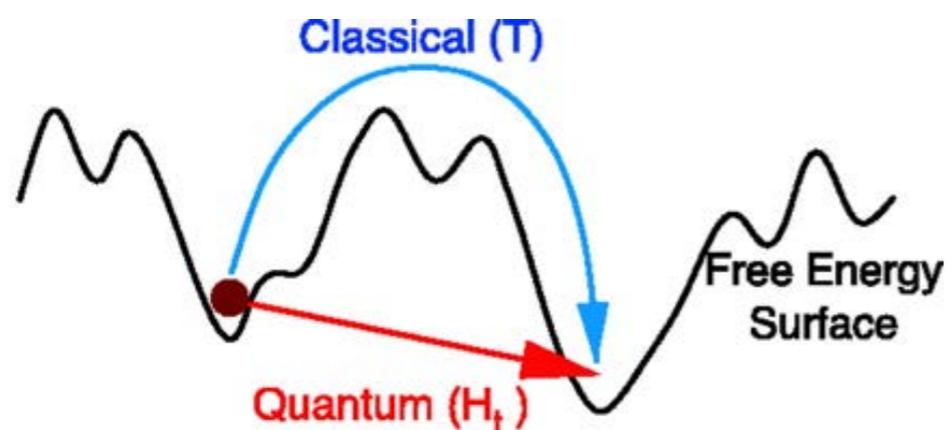
Scaling of annealing times in quantum annealing

- Quantum adiabatic theorem

$$t \propto \frac{1}{\Delta_{\min}^2}$$

- Open systems tunneling through a barrier

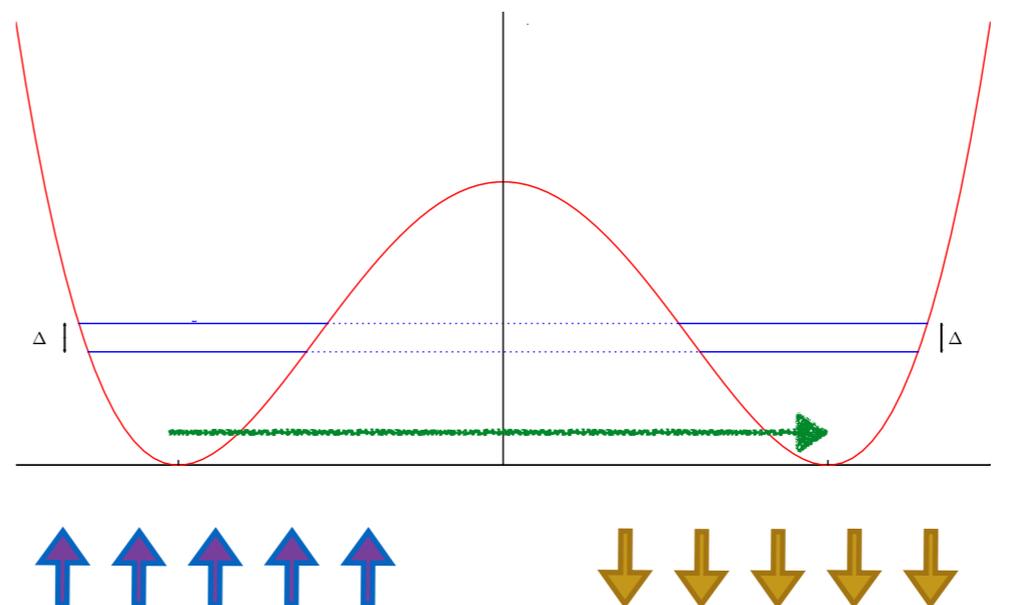
$$t \propto \epsilon^2 \propto \frac{1}{\Delta_{\min}^2}$$



S. Knysh, arXiv:1506.08608

Tunneling between two states of a ferromagnet

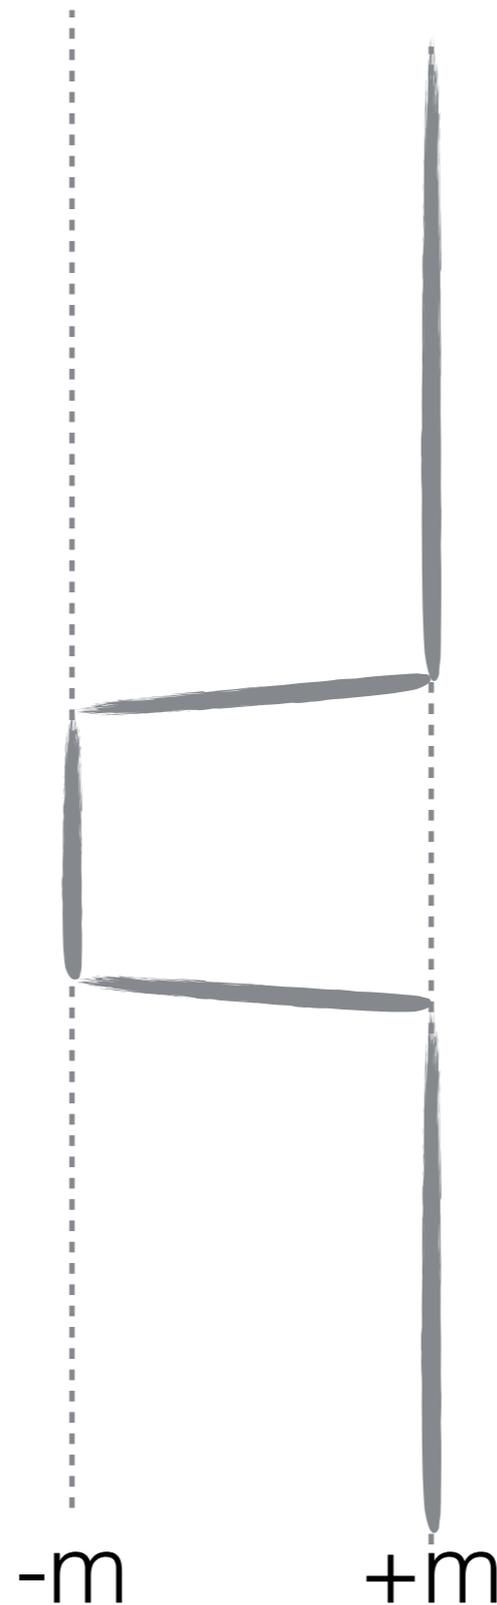
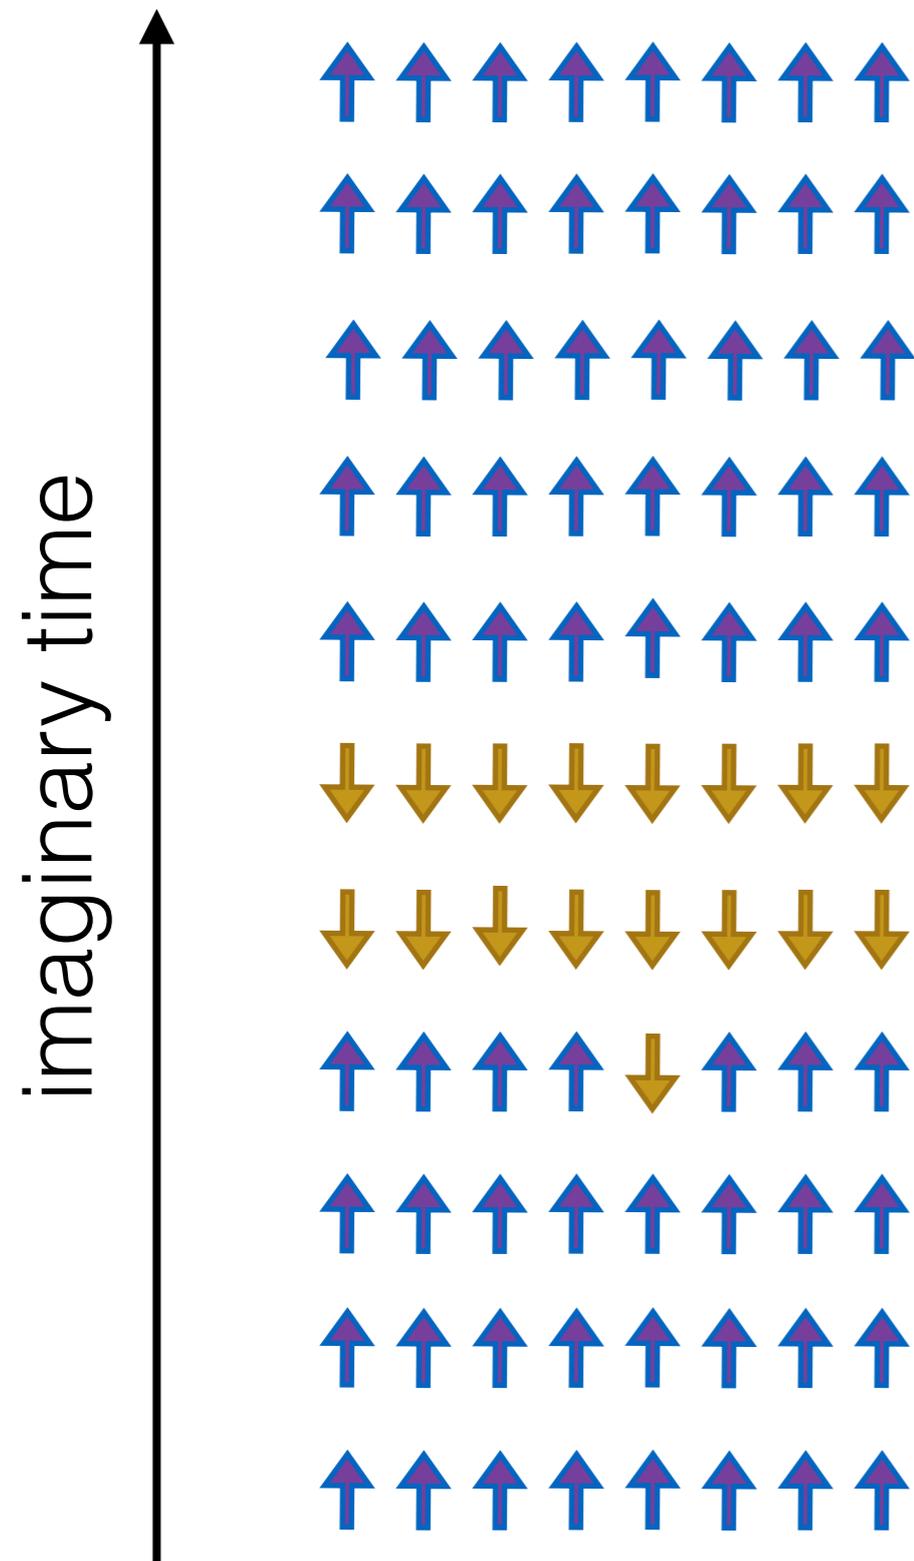
$$H = - \sum_{\langle ij \rangle} J \sigma_i^z \sigma_j^z - \Gamma \sum_i \sigma_i^x$$



energy splitting Δ

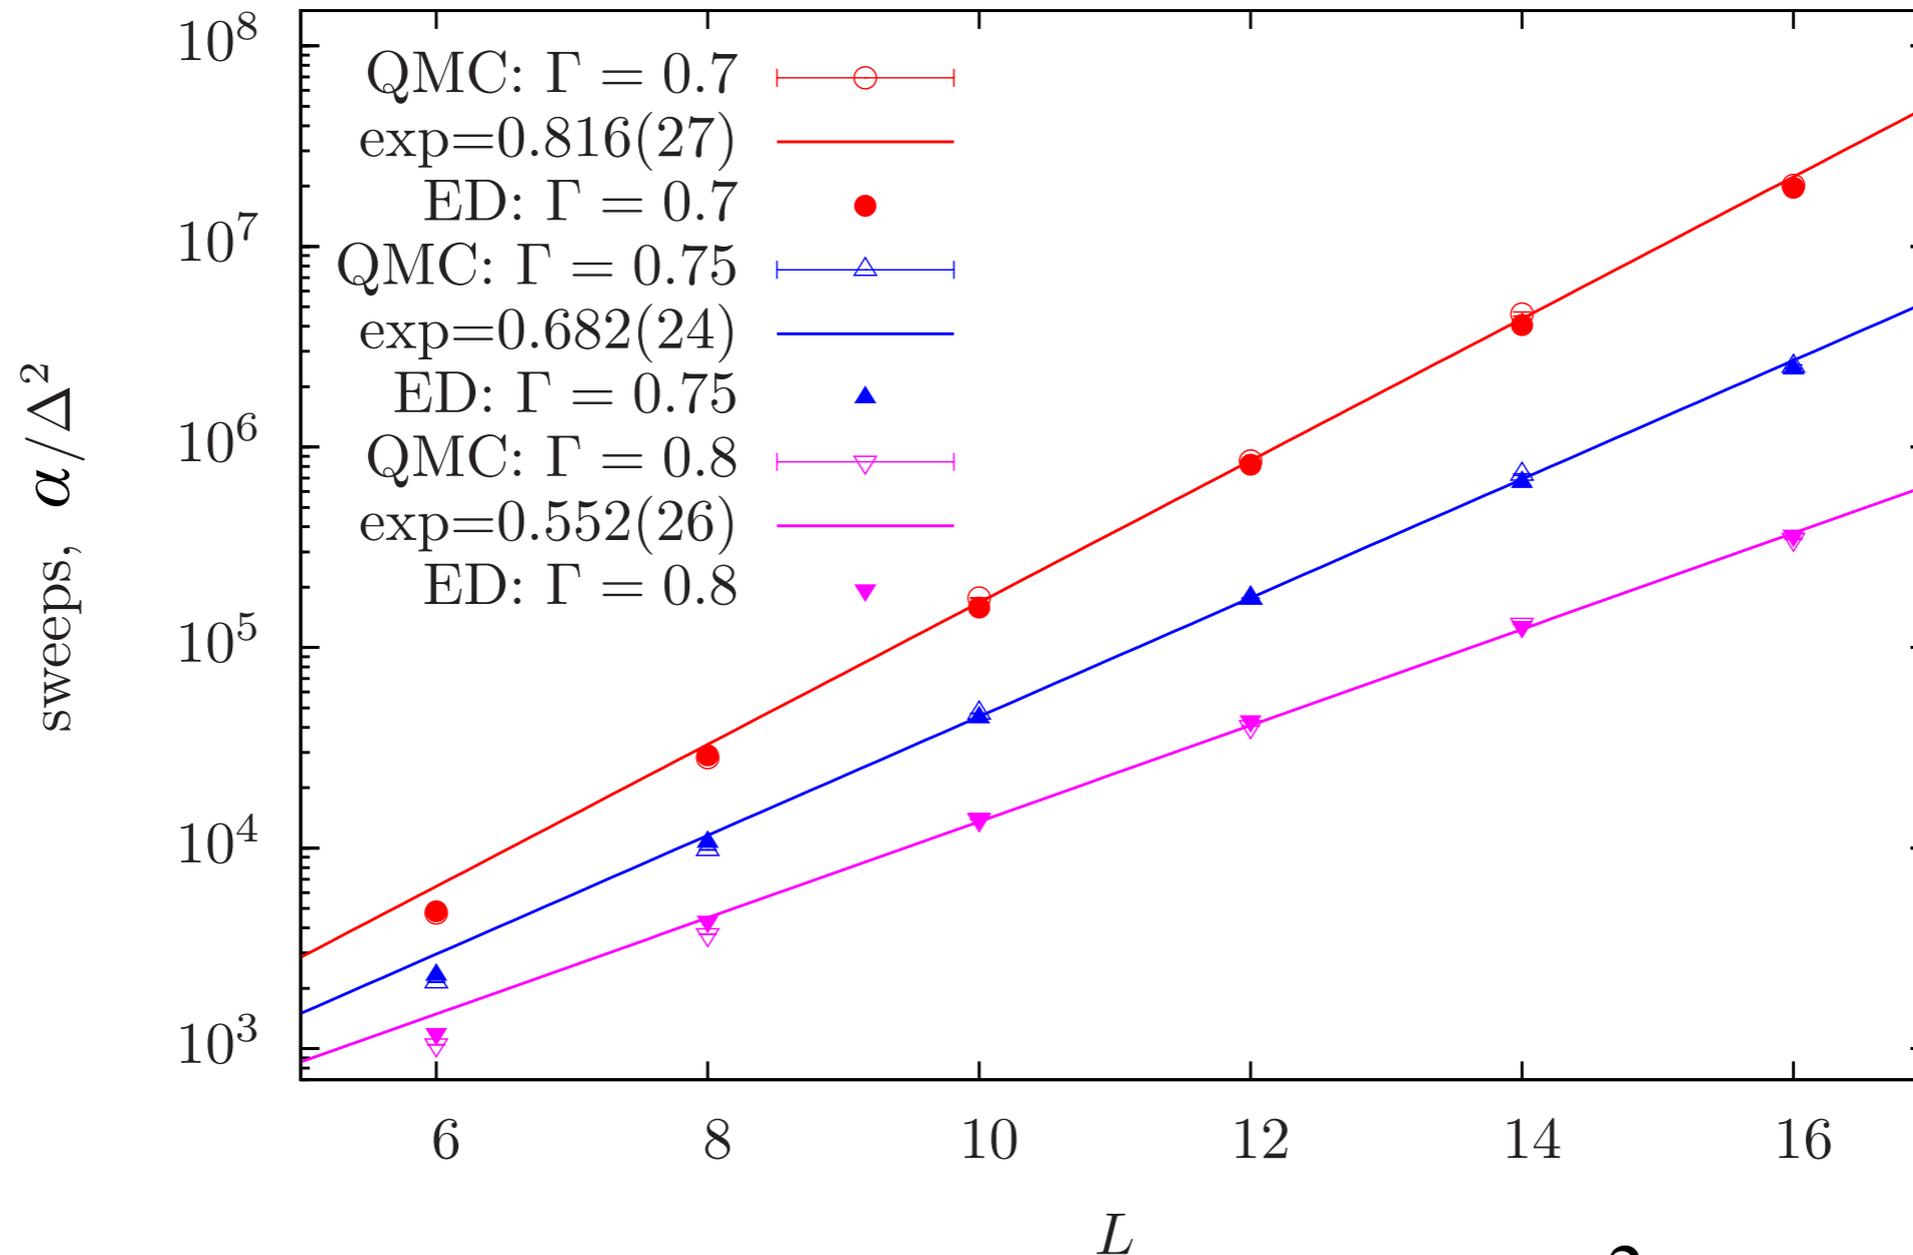
$$\psi_0 = \left| \begin{array}{c} \uparrow \uparrow \uparrow \uparrow \uparrow \\ \downarrow \downarrow \downarrow \downarrow \downarrow \end{array} \right\rangle + \left| \begin{array}{c} \downarrow \downarrow \downarrow \downarrow \downarrow \\ \uparrow \uparrow \uparrow \uparrow \uparrow \end{array} \right\rangle$$

$$\psi_1 = \left| \begin{array}{c} \uparrow \uparrow \uparrow \uparrow \uparrow \\ \downarrow \downarrow \downarrow \downarrow \downarrow \end{array} \right\rangle - \left| \begin{array}{c} \downarrow \downarrow \downarrow \downarrow \downarrow \\ \uparrow \uparrow \uparrow \uparrow \uparrow \end{array} \right\rangle$$



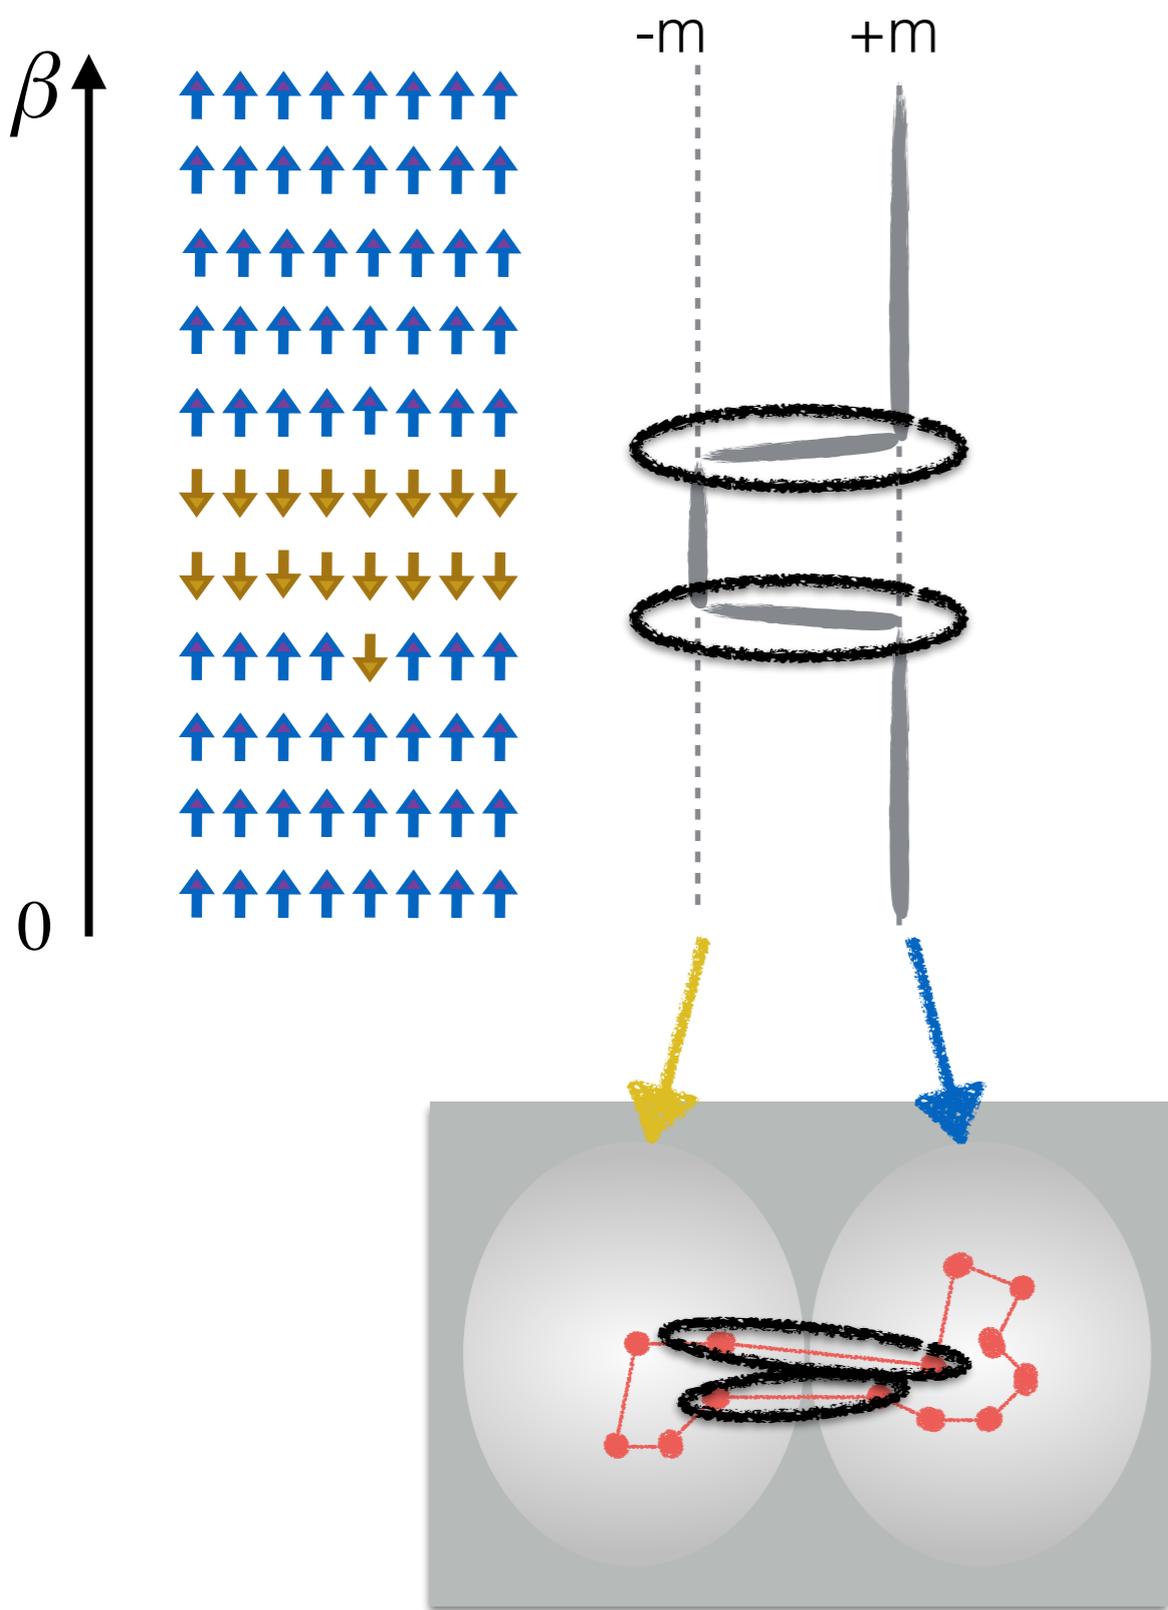
Tunneling determined from the time QMC needs to create an instanton.

Scaling with system size

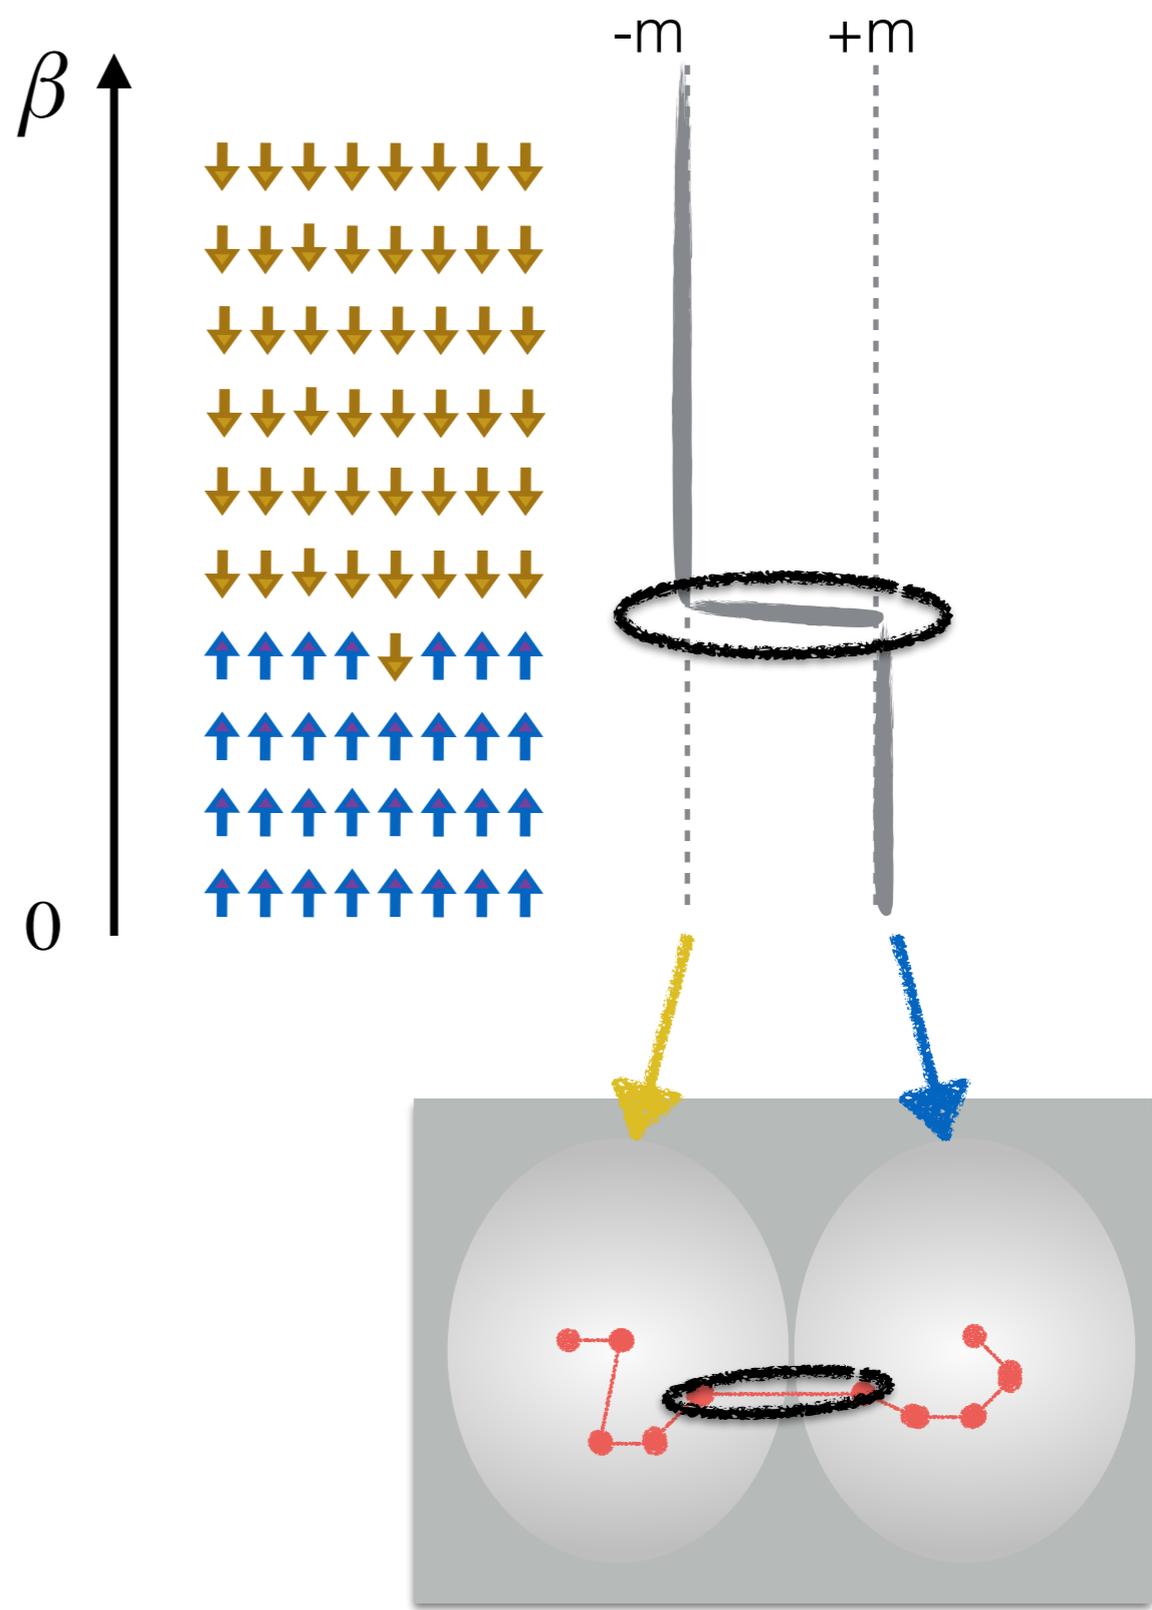


$T = J/16$

Tunneling rate given by QMC scale with $\sim \Delta^2$
as in physical dynamics

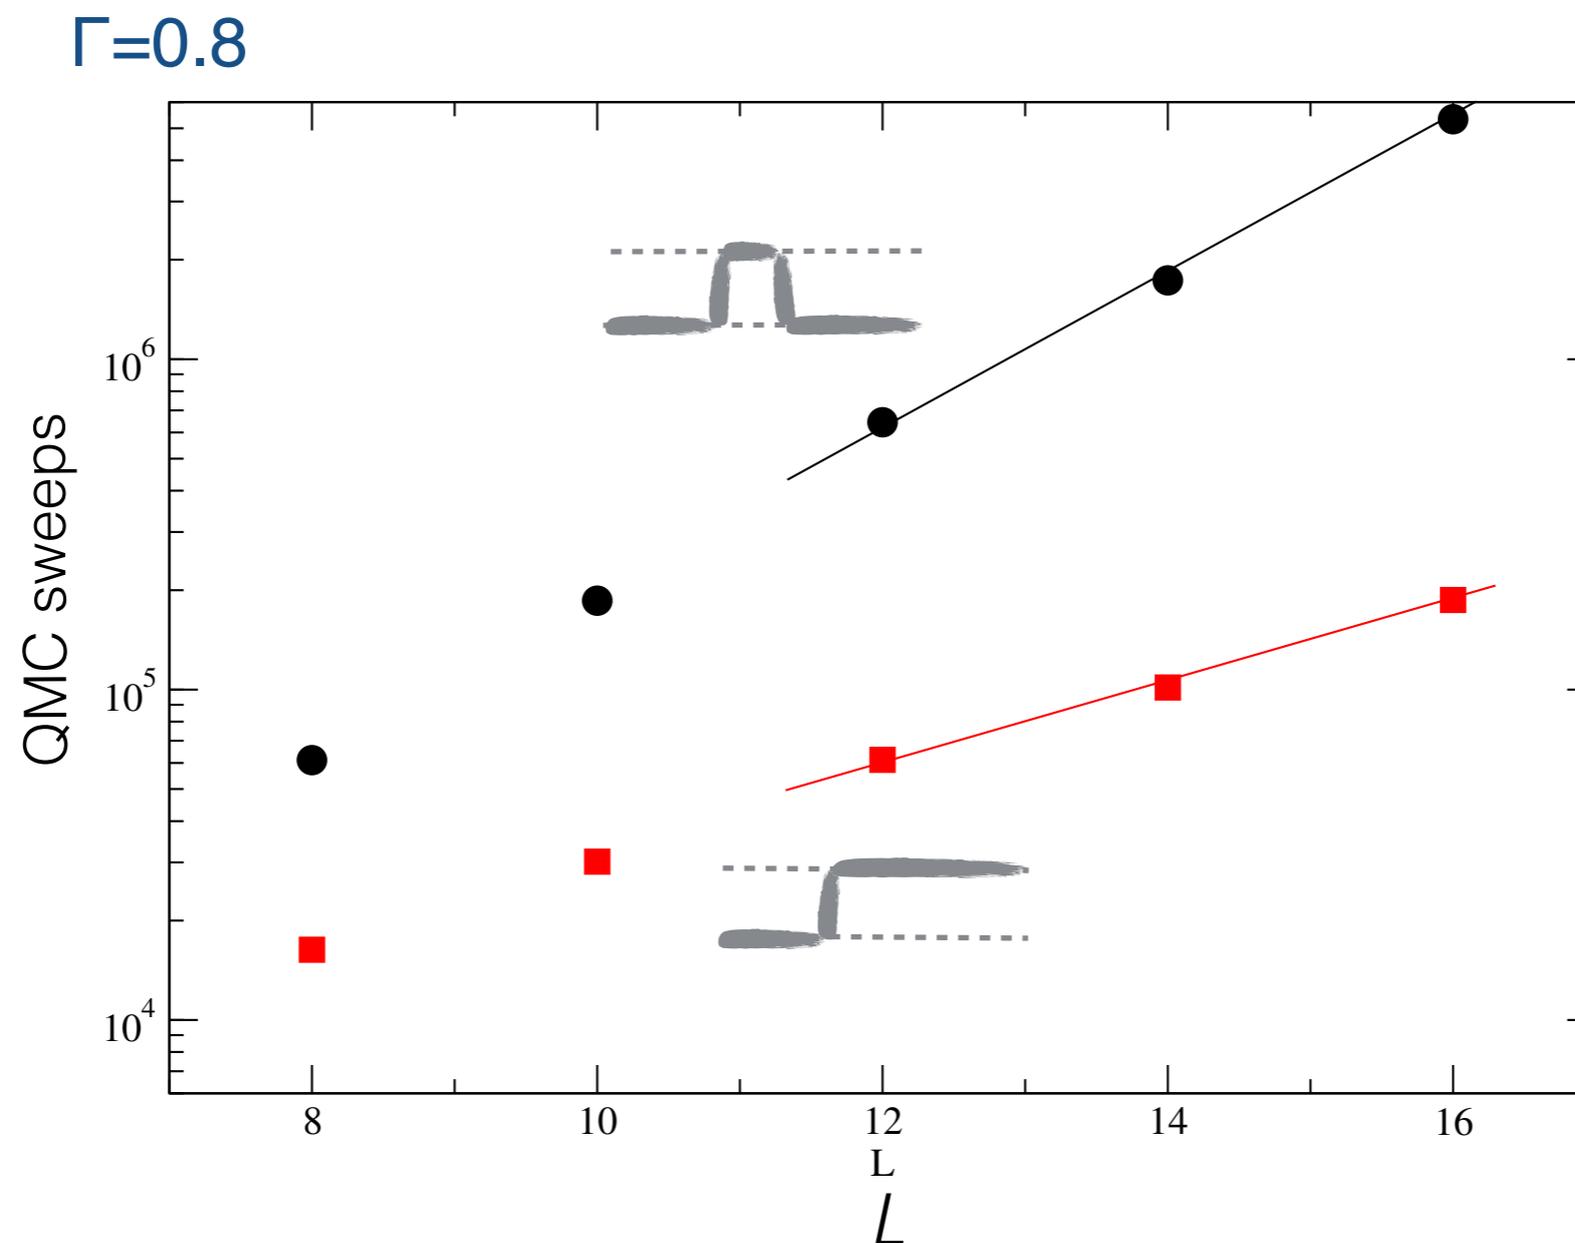


$$Z = \text{Tr} e^{-\beta H}$$



$$K(x_0; x_M) = \langle x_0 | e^{-\beta H} | x_M \rangle$$

Scaling with system size L



$$\tau_{PBC} \sim \frac{1}{\Delta^2}$$

$$\tau_{OBC} \sim \frac{1}{\Delta}$$

By cutting open the trace in imaginary time we have a new algorithm, with a better scaling.

Implications for quantum annealing

- QMC is often relevant for performance of a quantum annealer
 - Scaling is the same for Grover search type problems
 - Scaling is the same for tunneling through barrier in a ferromagnet
 - Scaling of QMC with open temporal boundary conditions has a quadratic speedup over QA for tunneling through barrier in a ferromagnet
- When can QA outperform QMC?
 - QA has a large constant advantage over QMC
 - non-stoquastic driver Hamiltonians don't have a QMC mapping and may be able to outperform stoquastic QA
 - QMC dynamics may have topological obstructions in some cases (Hastings and Freedman)

Summary

- Simulations have been performed for three types of problems
 - small tailored instances (clusters, chains, ...)
 - random instances engineered to enhance quantum enhancement potential
 - application problems
- Software framework for simulations
 - developed a framework to efficiently perform many millions of simulations
 - implemented a portfolio of simulation algorithms and SOA optimizers
 - built a library of hard problem instances
- QMC is a predictor for QA performance on a broad class of problems
 - QMC is seen to have same scaling as QA for tunneling through barrier
 - QMC as classical optimizer: use a large time step to cheat
 - QMC to mimic physical device: use continuous time or small time steps
- Beware of the tails
 - Very fat power-law tails in quantum annealing
 - Fast instead of adiabatic annealing is advantageous for hard problems