How "Quantum" is the D-Wave Machine?

Seung Woo Shin

UC Berkeley

Simons Institute 3/27/2014

(Joint work with Graeme Smith, John A. Smolin, and Umesh Vazirani)

D-Wave’s Year of Computing Dangerously

Commercial quantum computer leaves PC in the dust

D-Wave’s Quantum Computer Goes to the Races, Wins

Tests suggest t
technology can

SCIENCE

D-Wave is quantum say USC scientists

Sceptics remain wary

Study doubts quantum computer speed

New benchmarks raise doubt over D-Wave’s ‘quantum computer,’ but Google is optimistic long-term

D-Wave's Quantum Computing Claim Disputed Again

Quantum Research Shows D-Wave's Computers Are (Probably) the Real Deal

Jamie Condliffe
Filed to: QUANTUM COMPUTING Today 6:28am

kind of optimisation prob
equation by choosing th
the problem crops up in
recognition and machin

McGeoch and her colle
Burnaby, ran the proble
formed from supercondu

D-Wave's Year of Computing Dangerously

Commercial quantum computer leaves PC in the dust

D-Wave's Quantum Computer Goes to the Races, Wins

SCIENCE

D-Wave is quantum say USC scientists

Sceptics remain wary

Are we meaning the same thing, when we say "quantum"?

New benchmarks raise doubt over D-Wave's 'quantum computer,' but Google is optimistic long-term

D-Wave's Quantum Computing Claim Disputed Again

Quantum Research Shows D-Wave's Computers Are (Probably) the Real Deal

Jamie Condliffe

Filed to: QUANTUM COMPUTING Today 6:28am

What is a "quantum device"?

"quantumness"

Quantum physics is important in the design of transistors.

But at every useful level of abstraction, the laptop is of course classical . . .

What is a "quantum device"?

"quantumness"

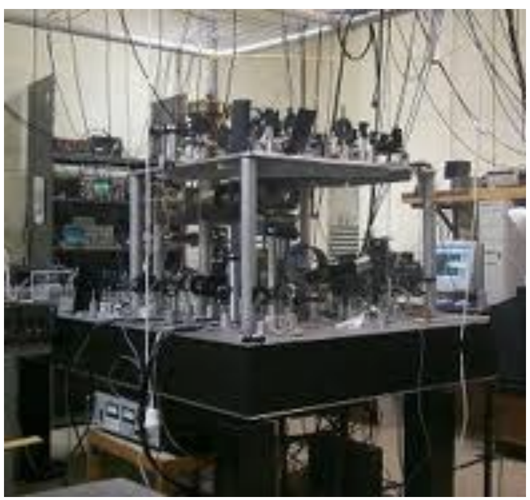

Minimal requirements for a QC

A) Large-scale quantum behavior
B) Suitable fault-tolerance
C) Universality
(or, demonstration of a useful algorithm)

What is a "quantum device"?

"quantumness"

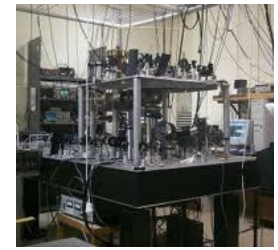

Minimal requirements for a QC

A) Large-scale quantum behavior
B) Suitable fault-tolerance
C) Universality
(or, demonstration of a useful algorithm)

1110101110101000 ⟶  ⟶ 0110101000101110

How do we study the physics of a black box?

Quantum tunneling? vs. Classical thermal effects?

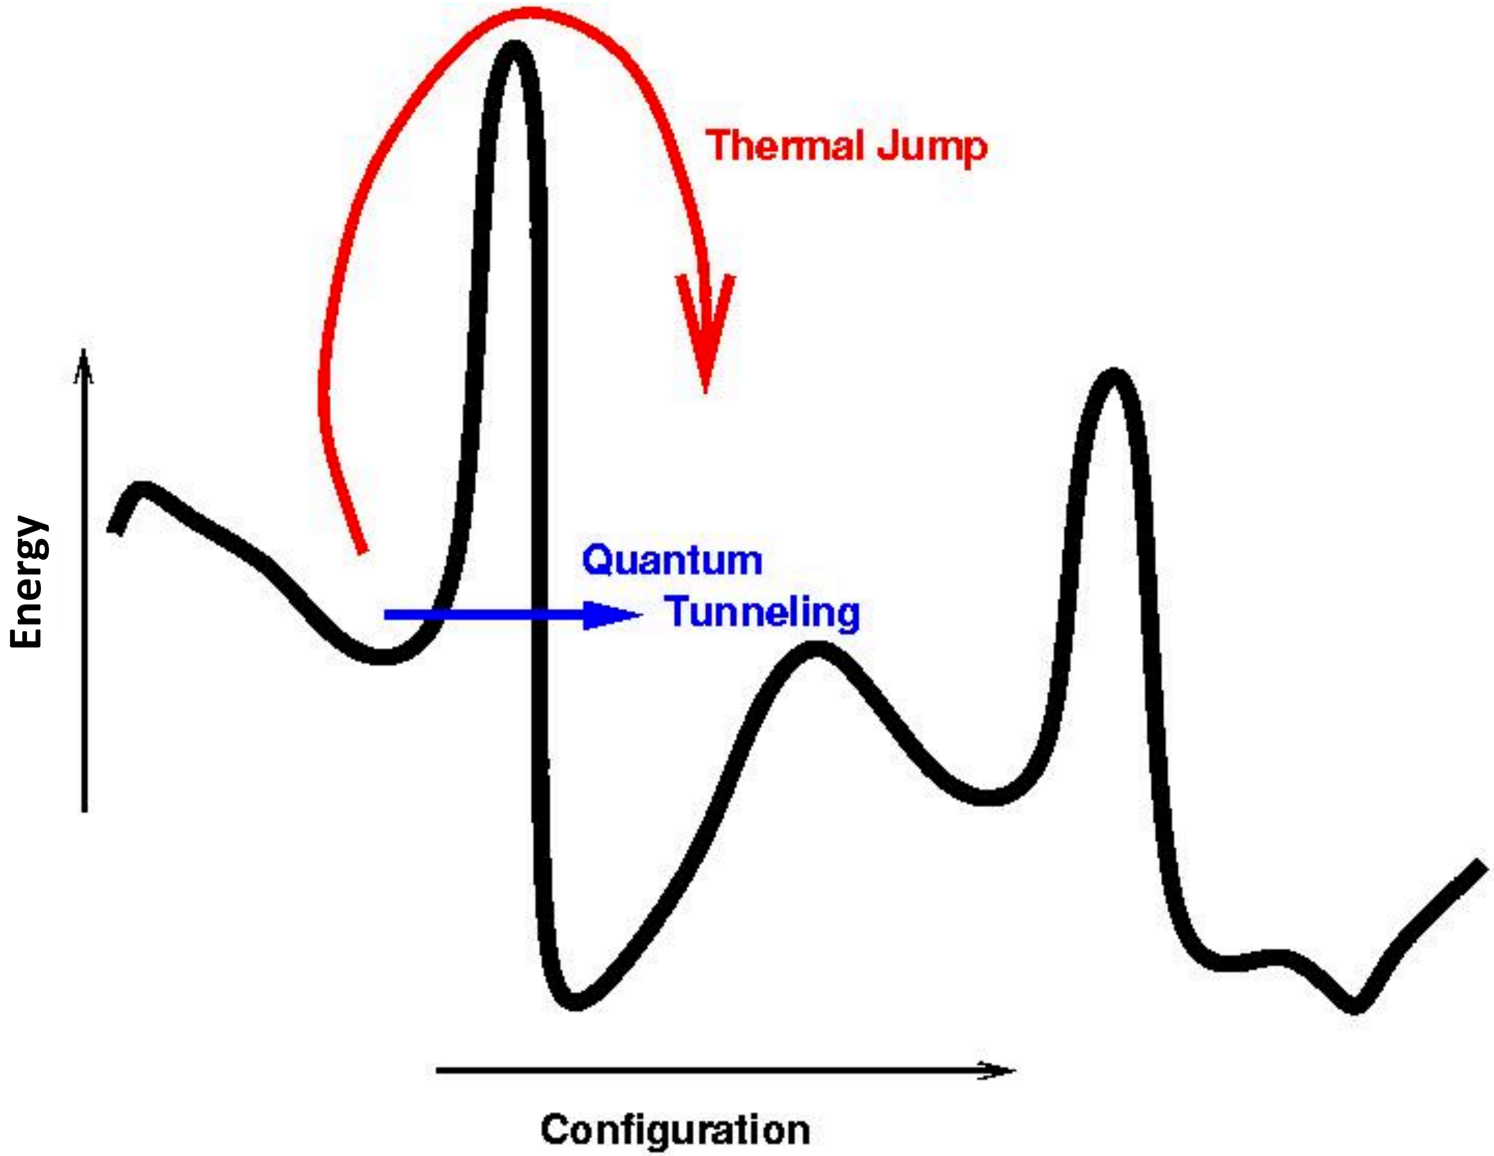

Image source: wiki/quantum_annealing

Boixo et al. 2013 arxiv:1304.4595

Compare D-Wave's input-output behavior to **classical simulated annealing** and **Quantum Monte Carlo** simulations!

1000 random instances, 1000 runs on each instance.

Success probabilities

	DW	SA	QMC
Instance #1	0.992	0.754	0.921
Instance #2	0.817	0.621	0.792
Instance #3	0.024	0.584	0.001
Instance #4	0.150	0.011	0.089
.	.	.	.
.	.	.	.
.	.	.	.
.	.	.	.
#1000	0.882	0.572	0.976

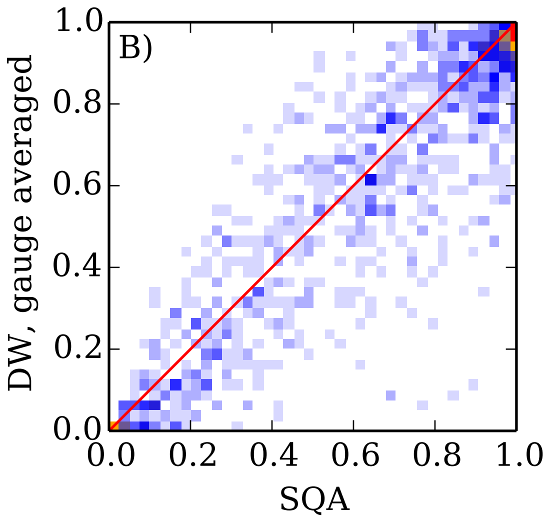

Quantum Monte Carlo

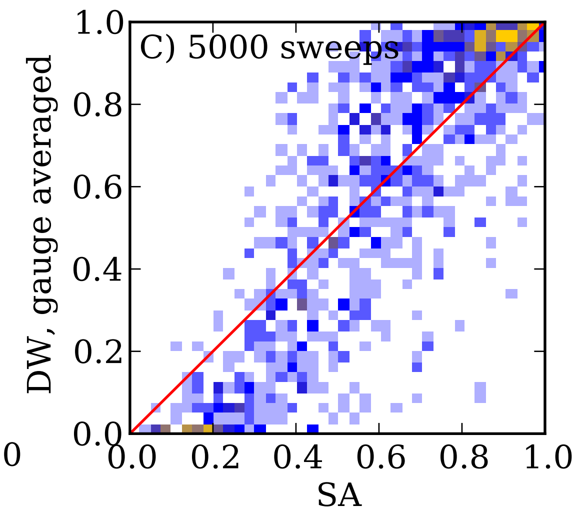

Simulated Annealing

Boixo et al. 2013 arxiv:1304.4595

Compare D-Wave's input-output behavior to **classical simulated annealing** and **Quantum Monte Carlo** simulations!

1000 random instances, 1000 runs on each instance.

Success probabilities

	DW	SA	QMC
Instance #1	0.992	0.754	0.921
Instance #2	0.817	0.621	0.792
Instance #3	0.024	0.584	0.001
Instance #4	0.150	0.011	0.089
.	.	.	.
.	.	.	.
.	.	.	.
.	.	.	.

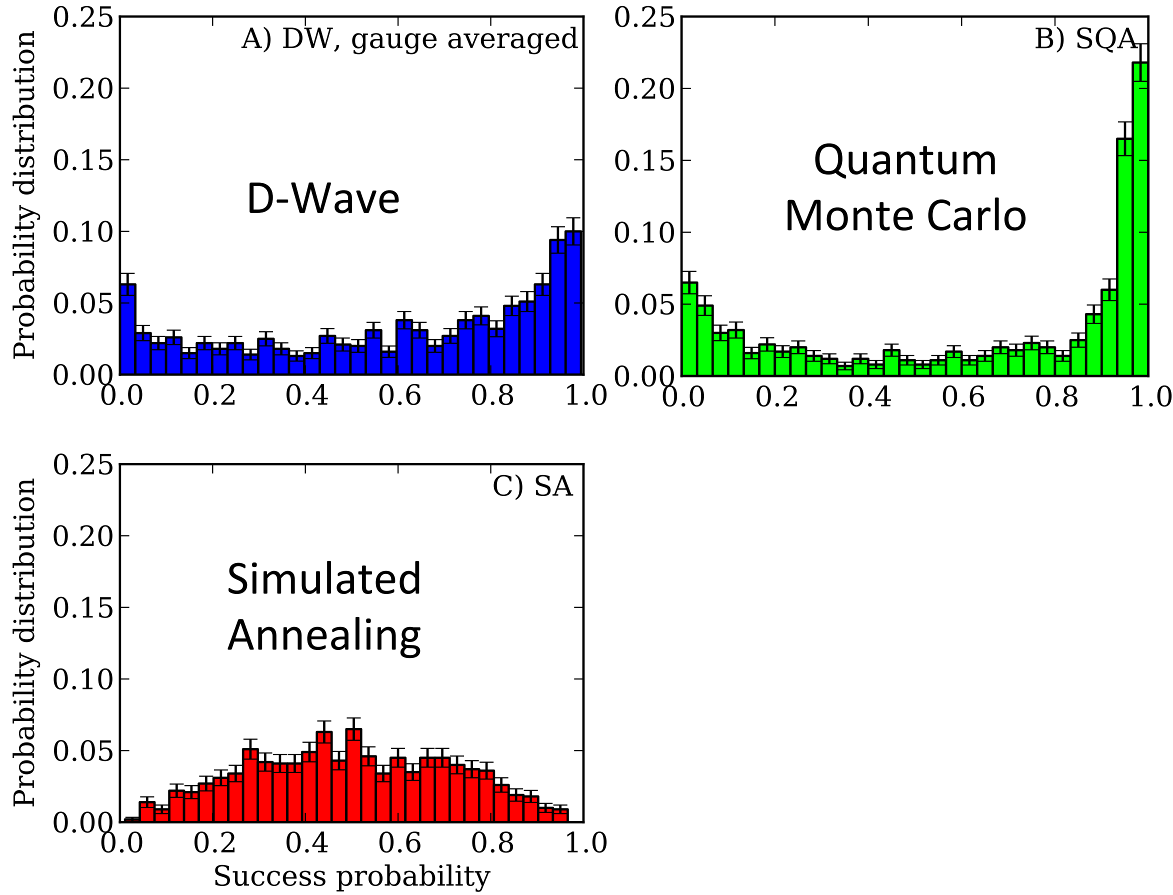

Quantum tunneling?

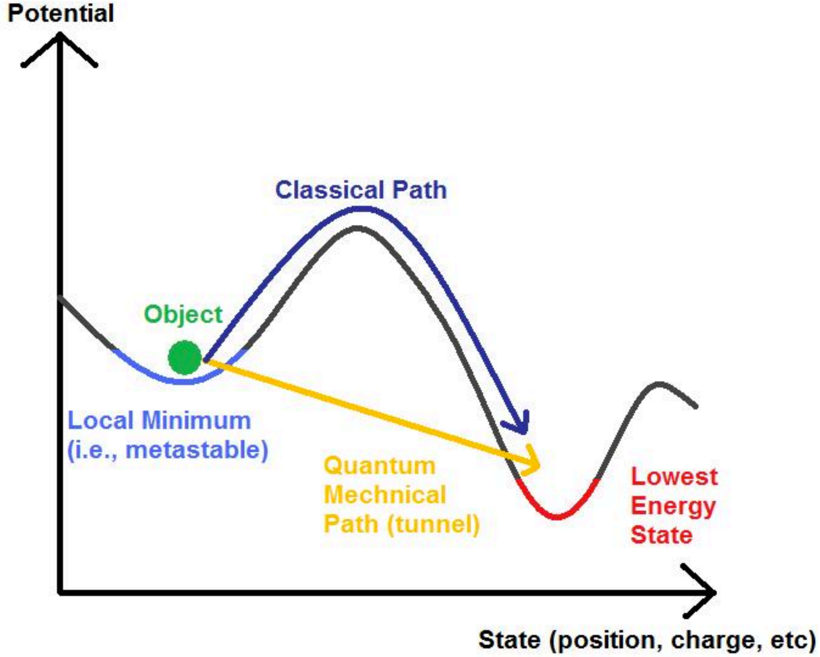

Poor correlation between the D-Wave machine and classical thermal annealing seemed to indicate that ***quantum tunneling*** was taking place in the machine.

Moreover, D-Wave seemed to work well on different sets of instances from classical thermal annealing.

Even if there is no speedup in general, maybe one could hope to identify a class of instances for which there will be speedup.

S, Smith, Smolin, Vazirani 2014 arxiv:1401.7087

A very simple *classical* model!

111010111010100 → 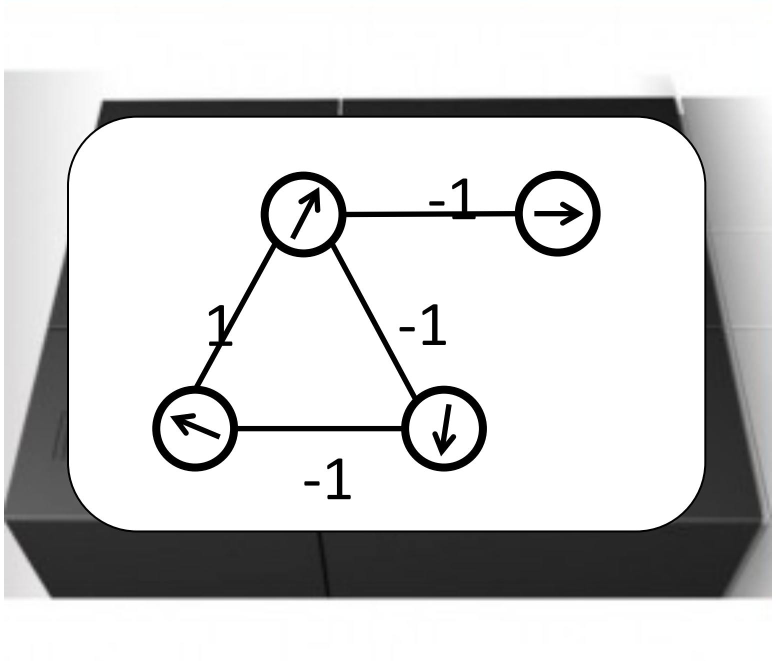 → 011010100010110

Using the same dataset of Boixo et al. (2013), we show that our classical model effectively describes D-Wave's algorithmic behavior.

A new classical model (S, Smith, Smolin, Vazirani 2014 arxiv:1401.7087)

Decoherence time of D-Wave's qubits is on the order of **nanoseconds**.
Computation time is on the order of **microseconds**.

$\Rightarrow$ We consider the ***simplest*** classical model that naturally arises from assuming that qubits decohere immediately. (Mean-field model)

Edges are Z-Z couplings

Nodes are superconducting flux qubits

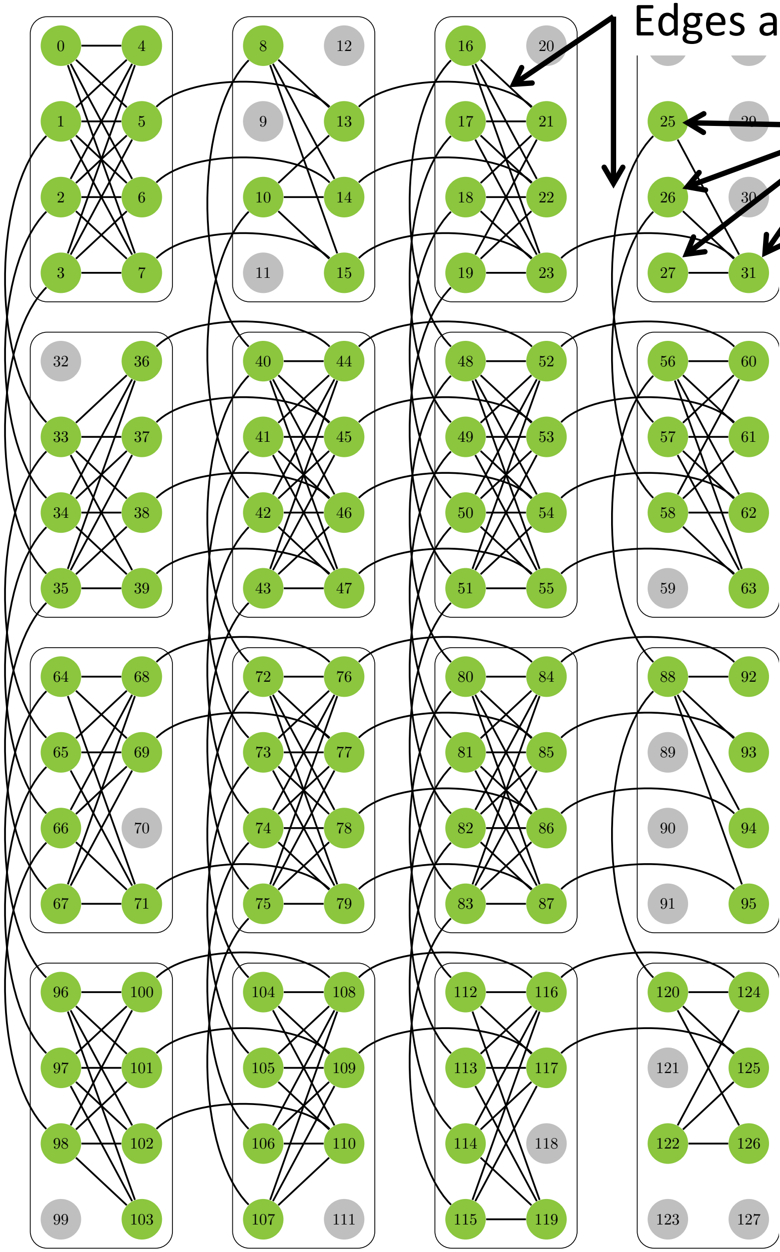

Figure: Boixo et al. (2013)

"Quantum Annealing"

$$H_0 = -\sum_i \sigma_i^x \qquad H_f = -\sum_{i\sim j} J_{ij}\sigma_i^z\sigma_j^z$$

$$H(t) = A(t)H_0 + B(t)H_f$$

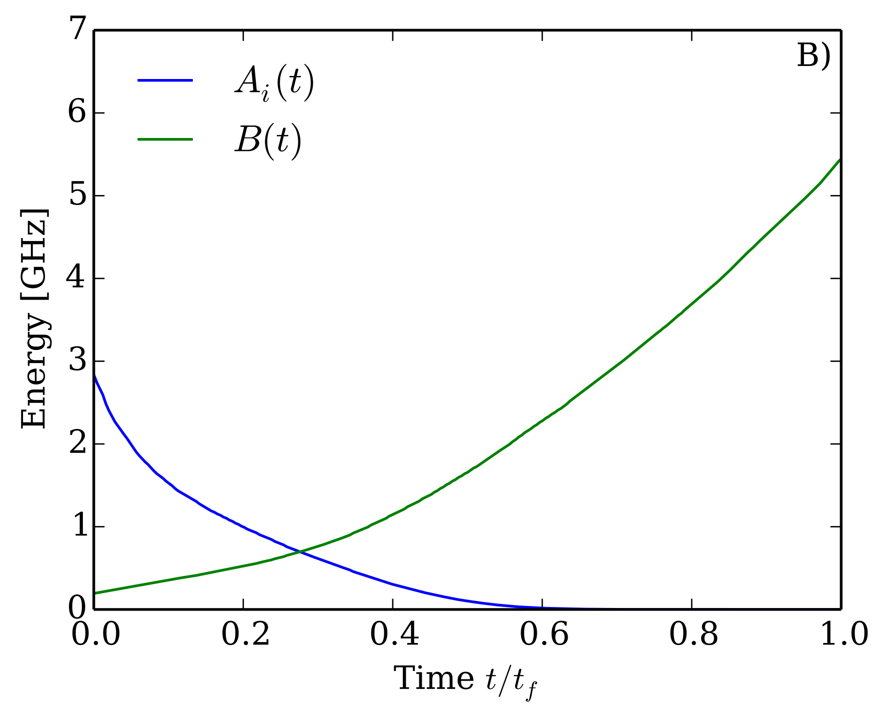

System temperature T>0 is constant and finite.

A new classical model **(S, Smith, Smolin, Vazirani 2014** arxiv:1401.7087**)**

Quantum Annealing

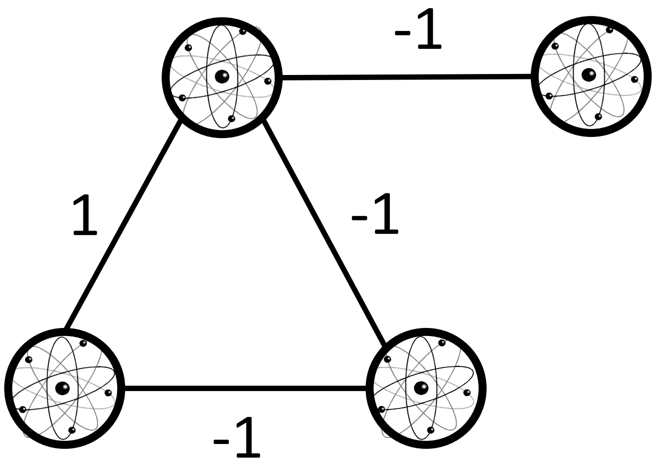

$$H_0 = -\sum_i \sigma_i^x$$

$$H_f = -\sum_{i \sim j} J_{ij} \sigma_i^z \sigma_j^z$$

Our classical model

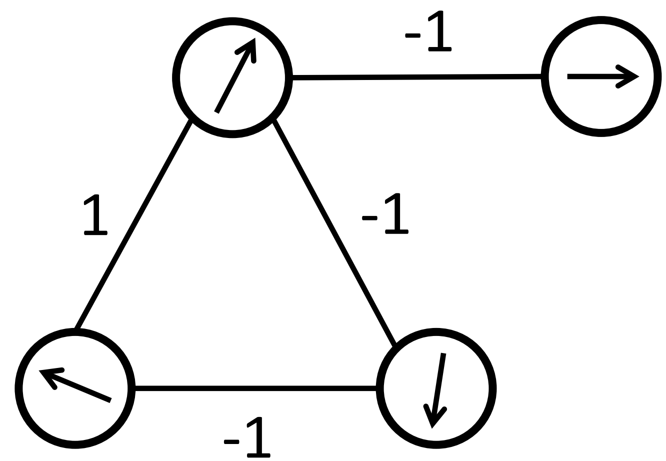

Model each spin by a classical magnet that points in some direction on the XZ plane

$$H_0 = -\sum_i \sin\theta_i$$

$$H_f = -\sum_{i \sim j} J_{ij} \cos\theta_i \cos\theta_j$$

A new classical model (S, Smith, Smolin, Vazirani 2014 arxiv:1401.7087)

$$H(t) = -A(t)\sum_i \sin\theta_i - B(t)\sum_{i\sim j} J_{ij}\cos\theta_i\cos\theta_j$$

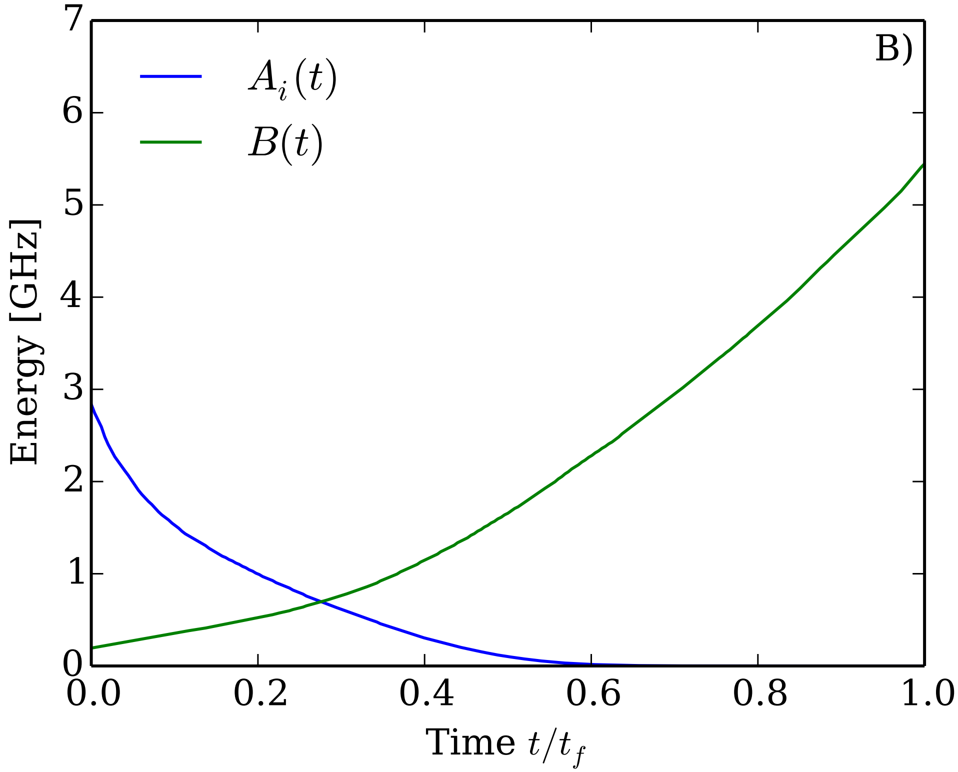

Temperature T>0 is constant.

Figure from Boixo et al.

A new classical model (S, Smith, Smolin, Vazirani 2014 arxiv:1401.7087**)**

$$H(t) = -A(t)\sum_i \sin\theta_i - B(t)\sum_{i\sim j} J_{ij}\cos\theta_i\cos\theta_j$$

Simulate using Metropolis algorithm

At each time step,

1. For each spin i, pick a new angle θ_i' at random.
2. Update spin i's state to θ_i' with probability $e^{-\Delta E_i/T}$.

N = 150,000 steps, T=0.22

A new classical model **(S, Smith, Smolin, Vazirani 2014** arxiv:1401.7087**)**

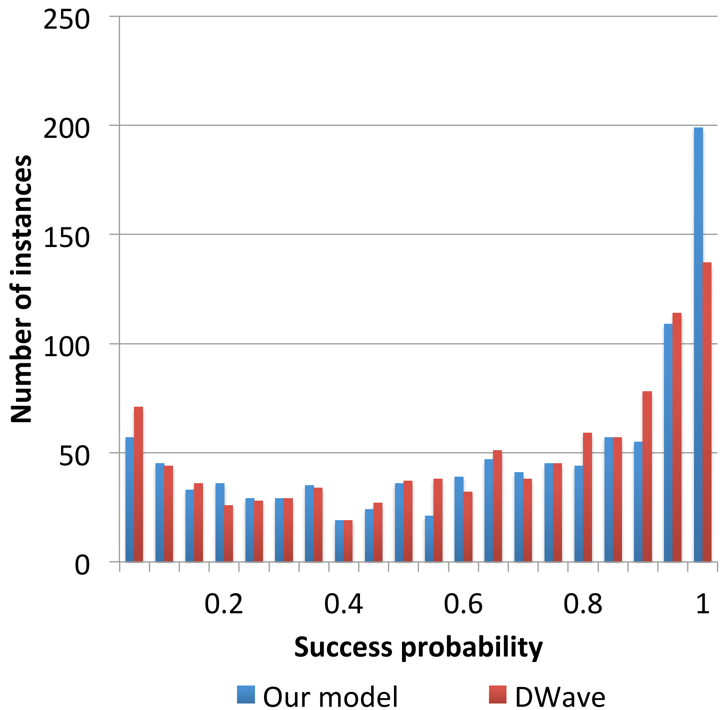

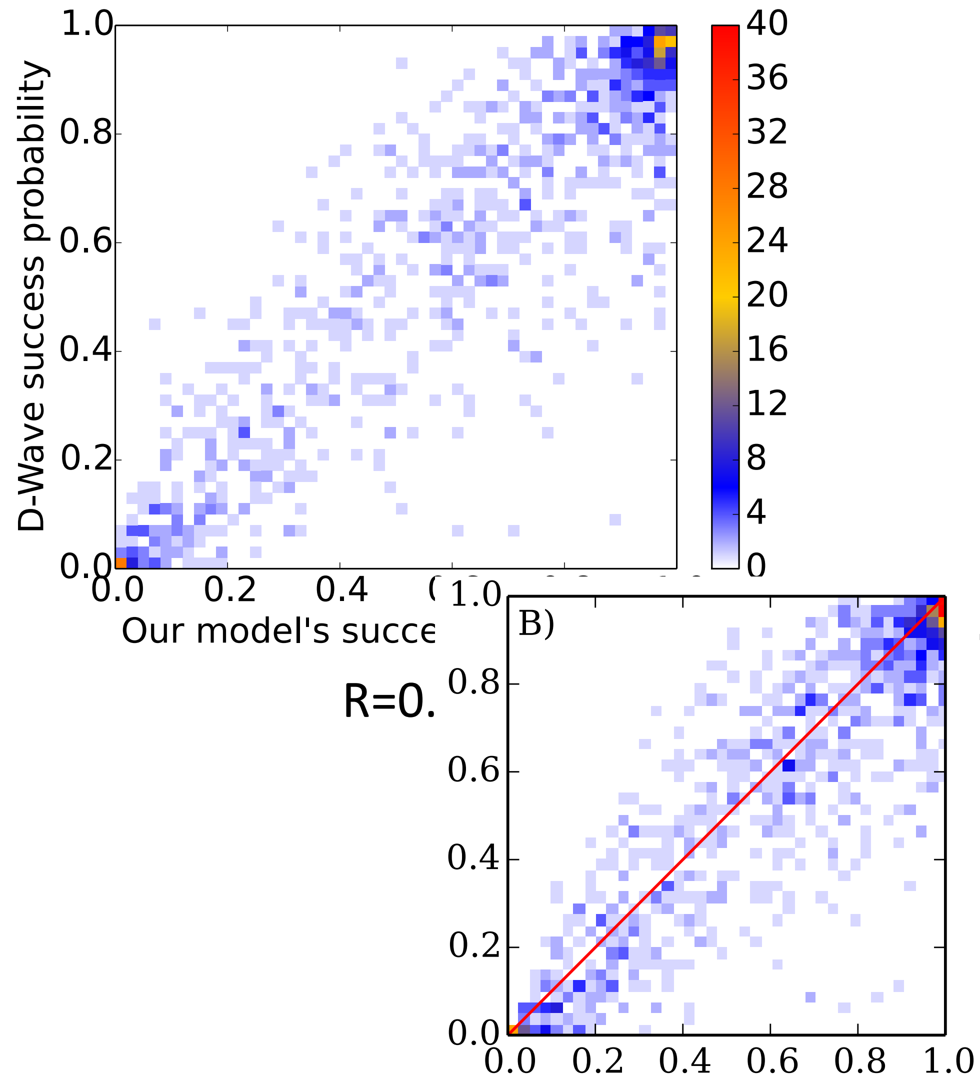

Algorithmic insights?

Why and how our model reproduces what was thought to be **quantum tunneling**

Our model – transverse field = simulated annealing

$$H(t,\vec{\theta}) = \cancel{A(t)H_i(\vec{\theta})} + B(t)H_f(\vec{\theta})$$

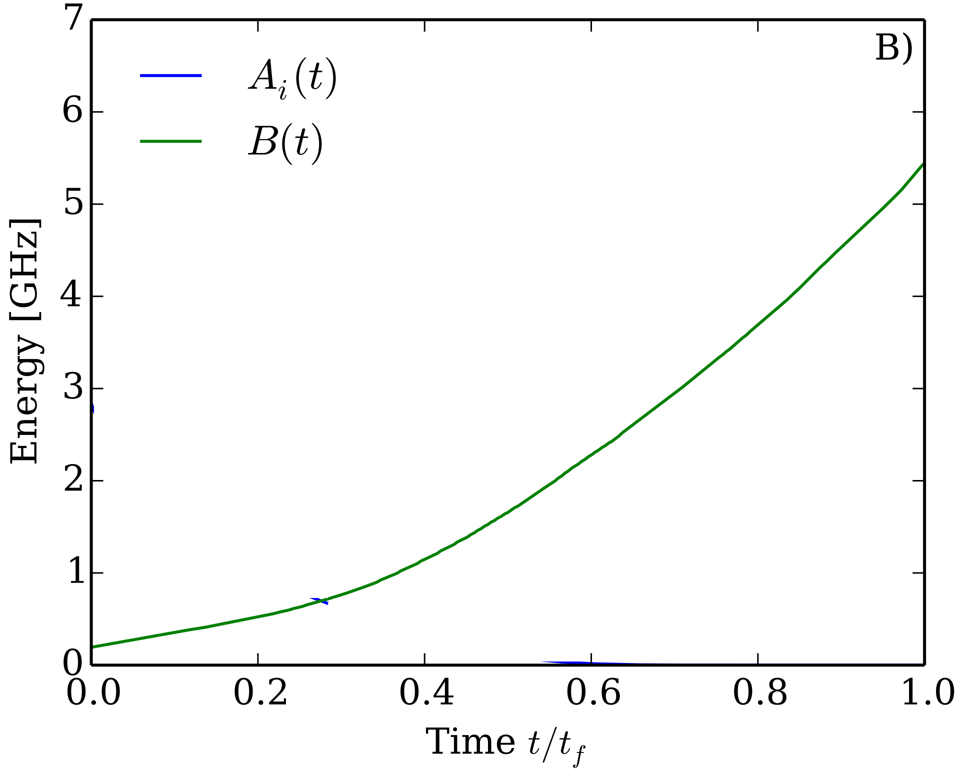

$$e^{-\frac{B(t)}{T}(H_f(\vec{\theta}_{\mathrm{new}})-H_f(\vec{\theta}_{\mathrm{old}}))}$$

Then, what is the role of transverse field?

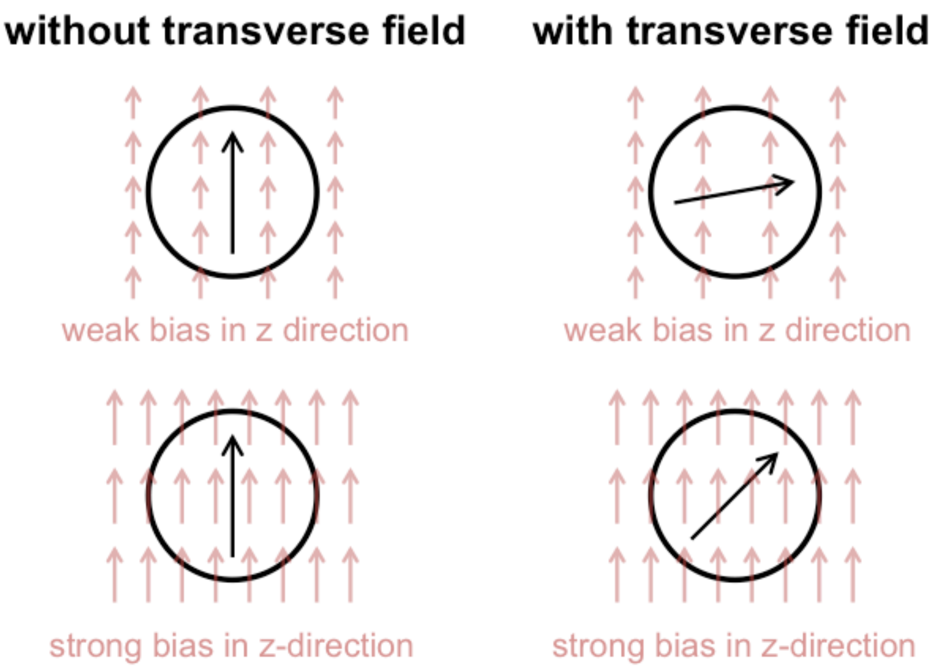

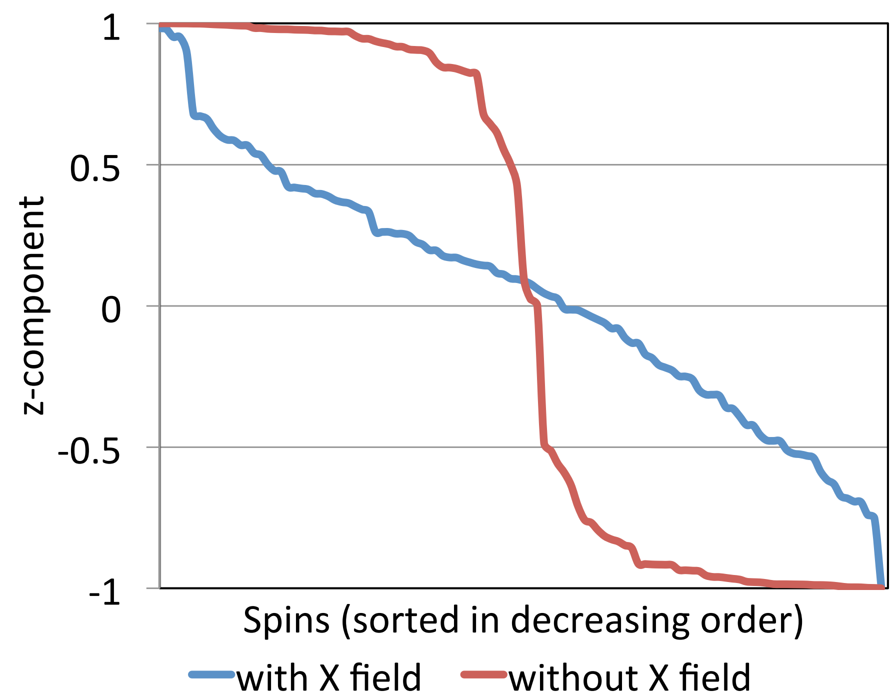

Reminiscent of the relationship between cuts and eigenvectors in spectral graph theory.

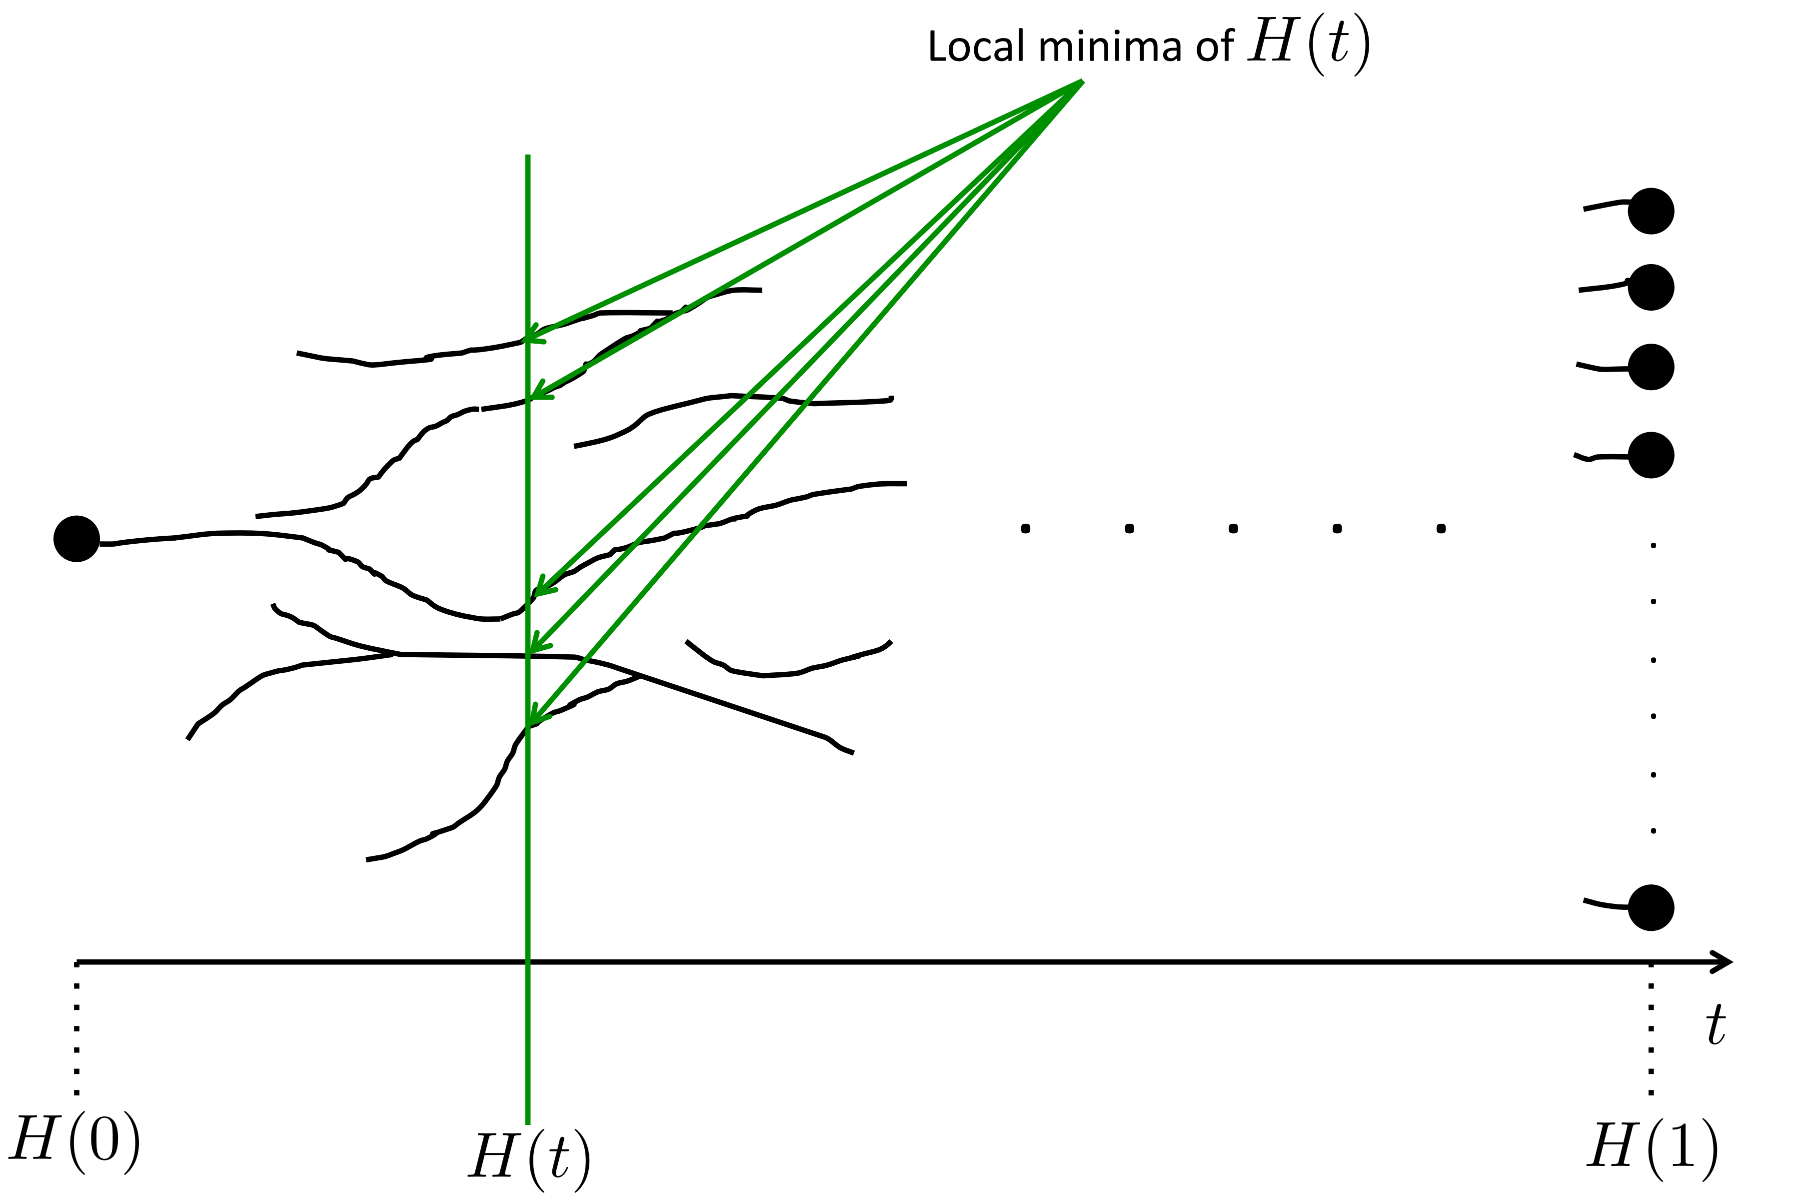
Local minima of $H(t)$
t
$H(0)$
$H(t)$
$H(1)$

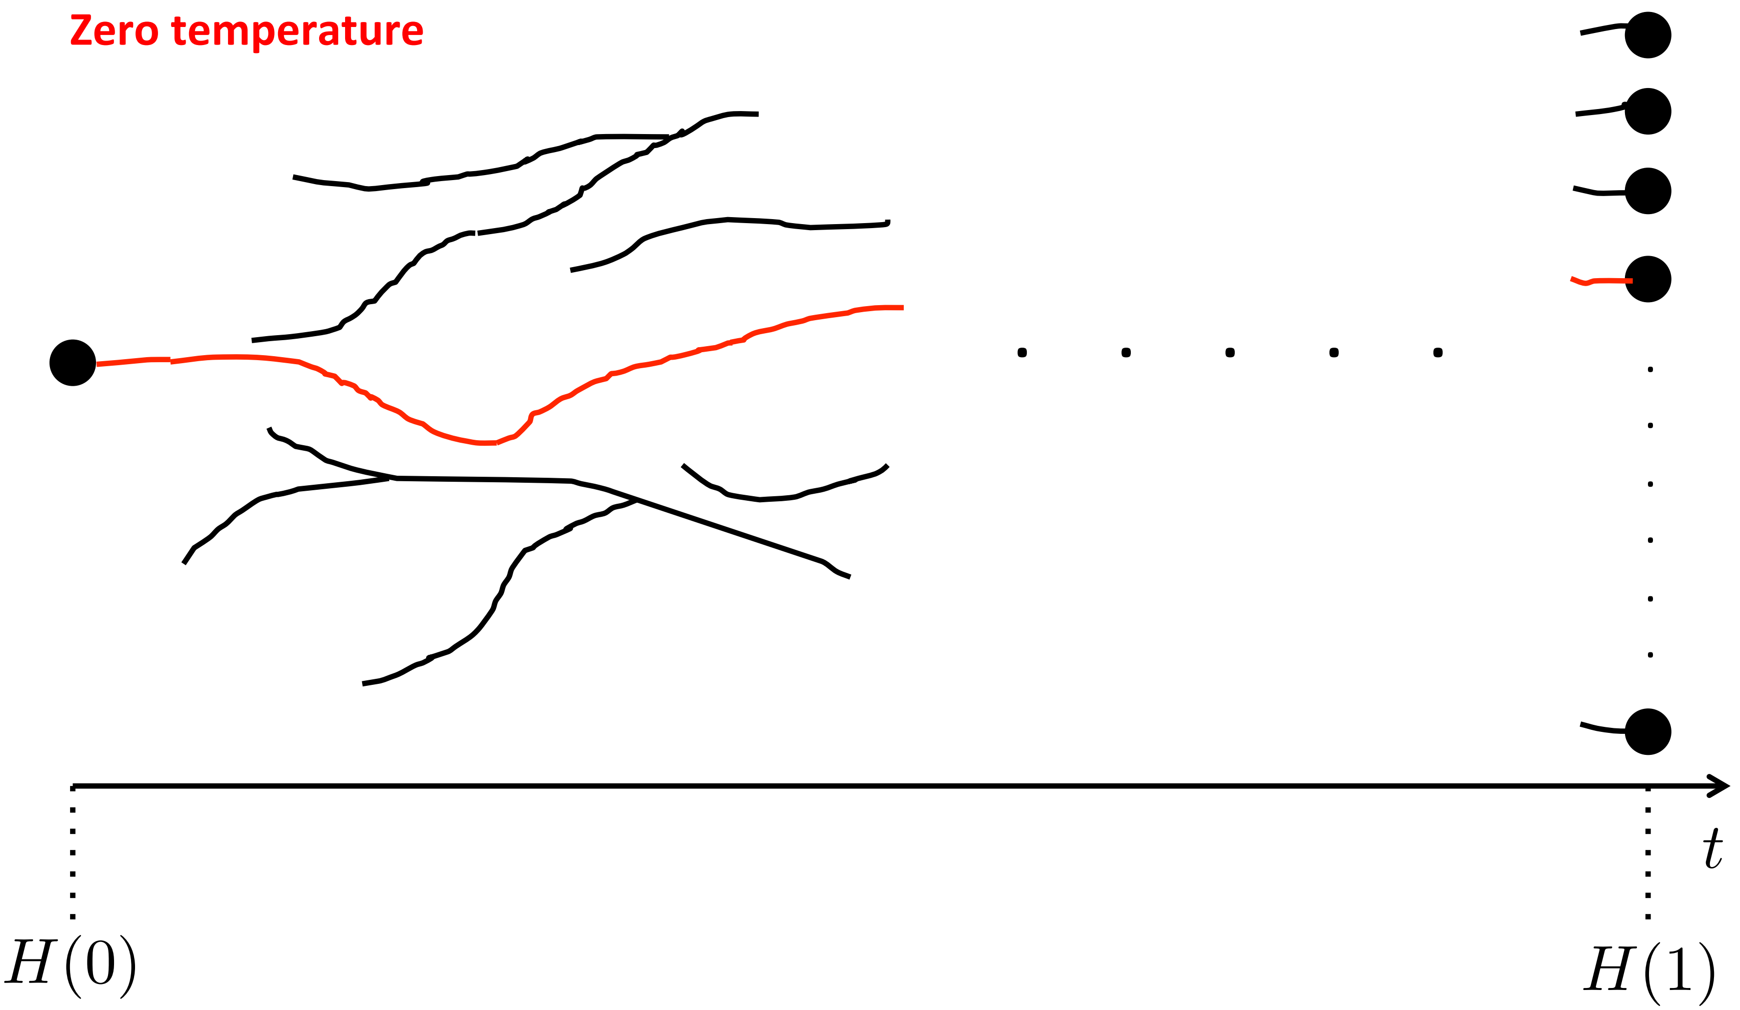

Zero temperature
t
$H(0)$
$H(1)$

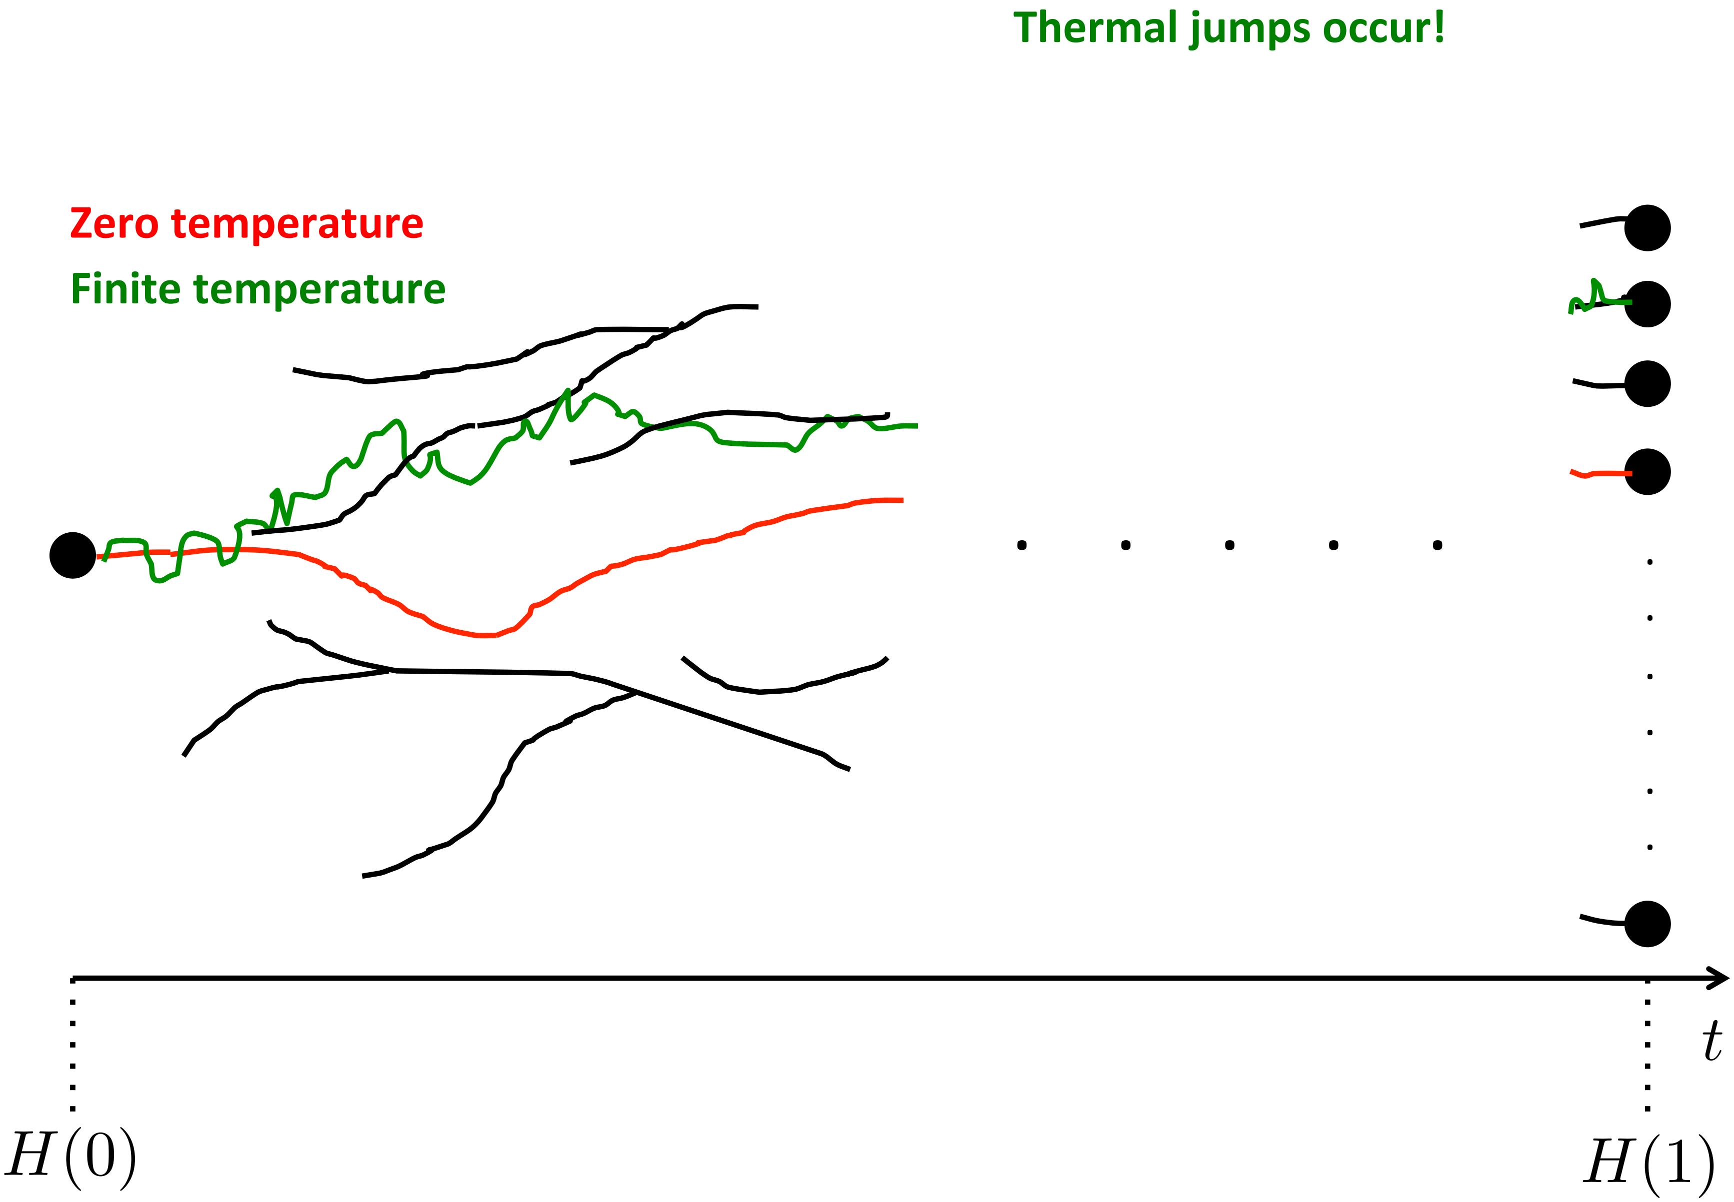

Thermal jumps occur!
Zero temperature
Finite temperature
t
$H(0)$
$H(1)$

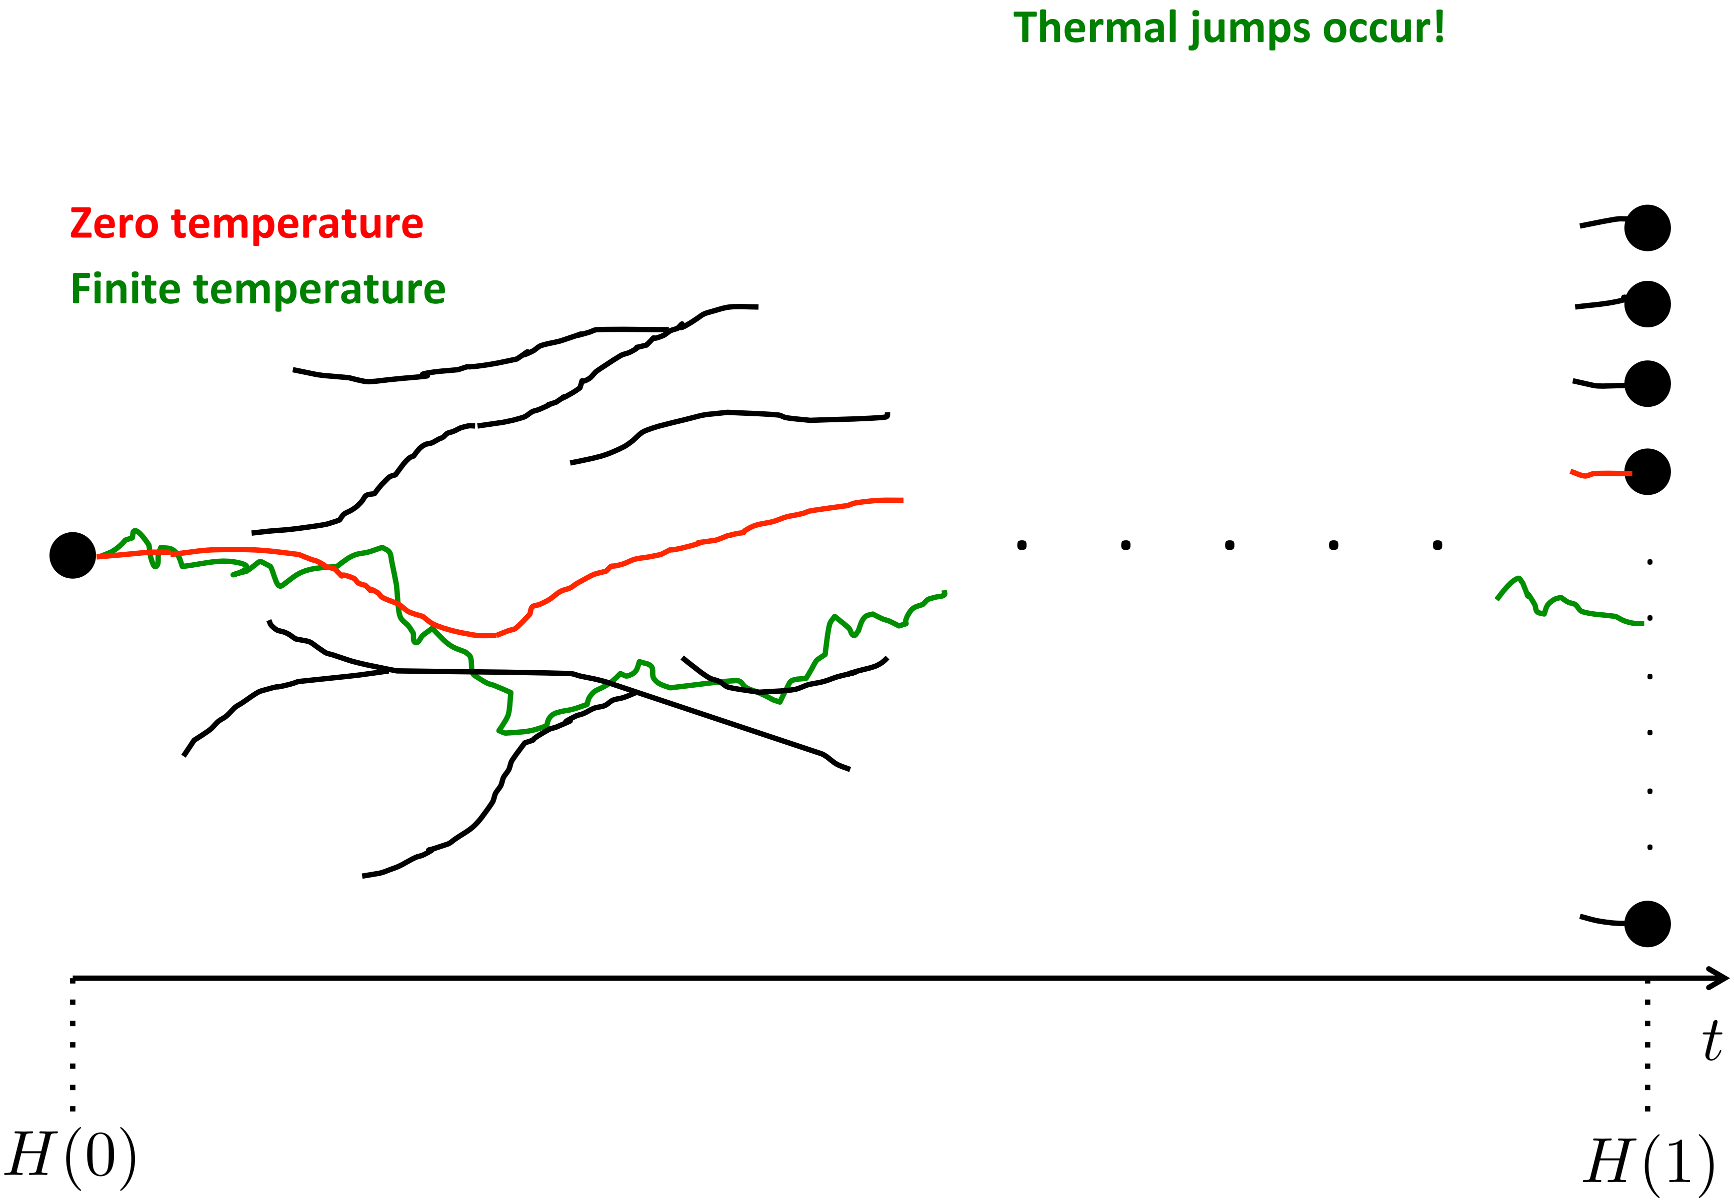
Thermal jumps occur!
Zero temperature
Finite temperature
t
$H(0)$
$H(1)$

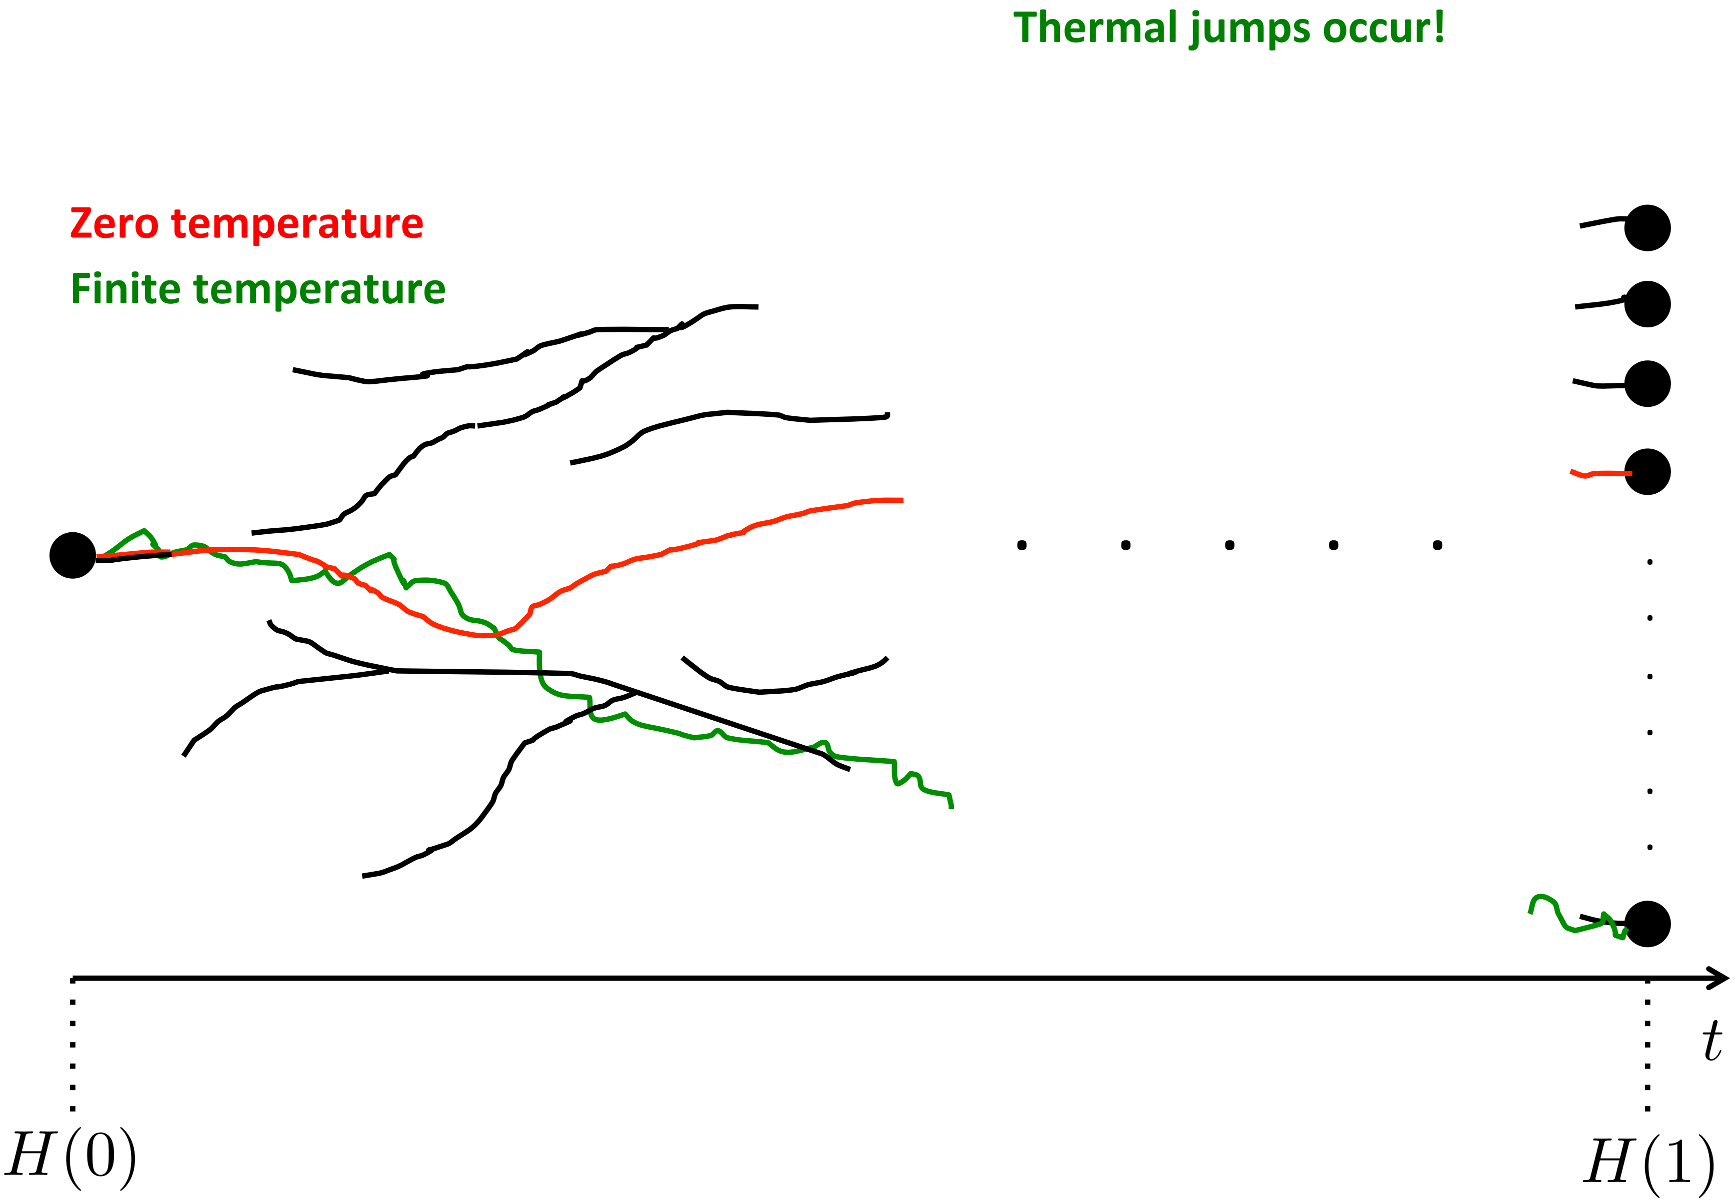

Thermal jumps occur!
Zero temperature
Finite temperature
t
$H(0)$
$H(1)$

Furthermore…

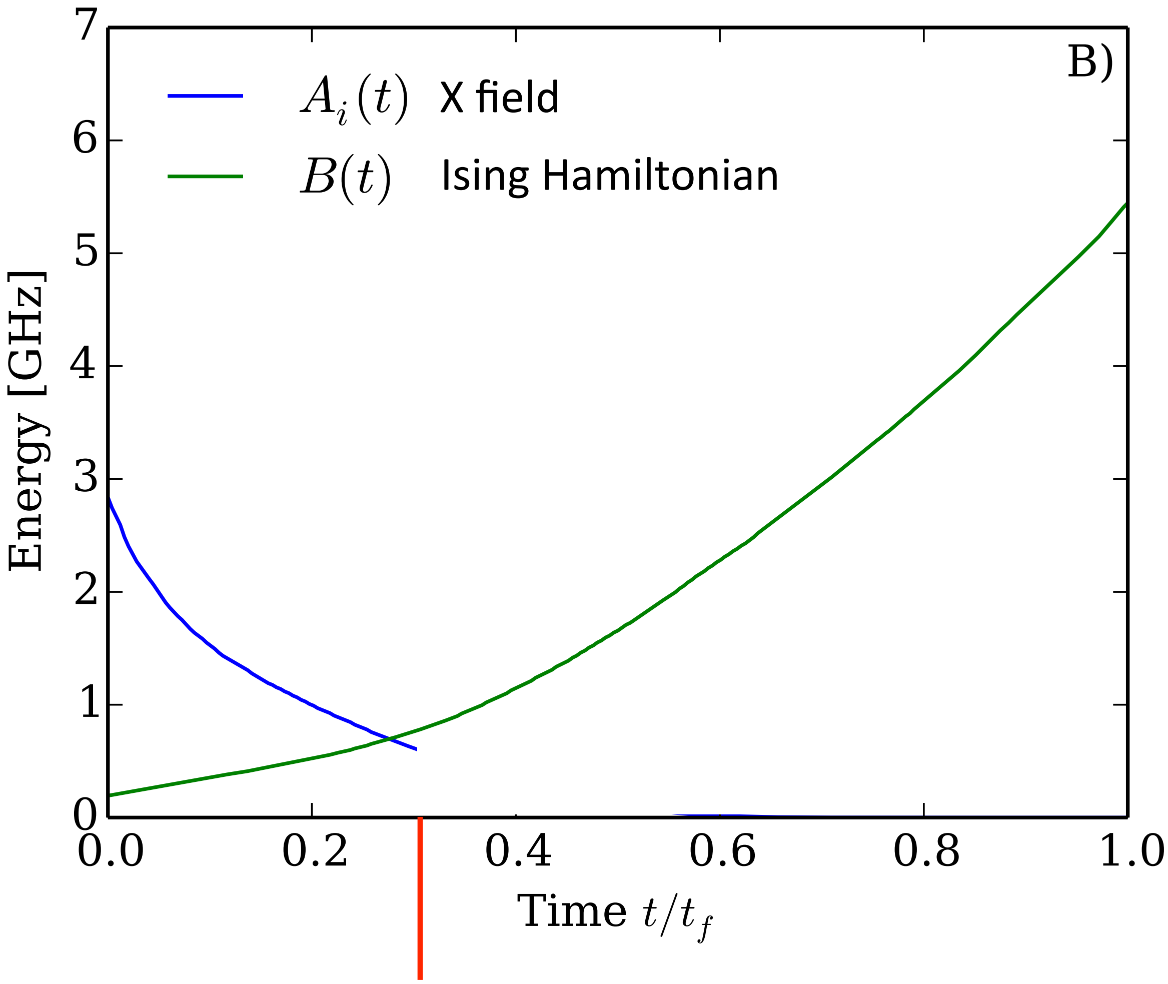

Furthermore…

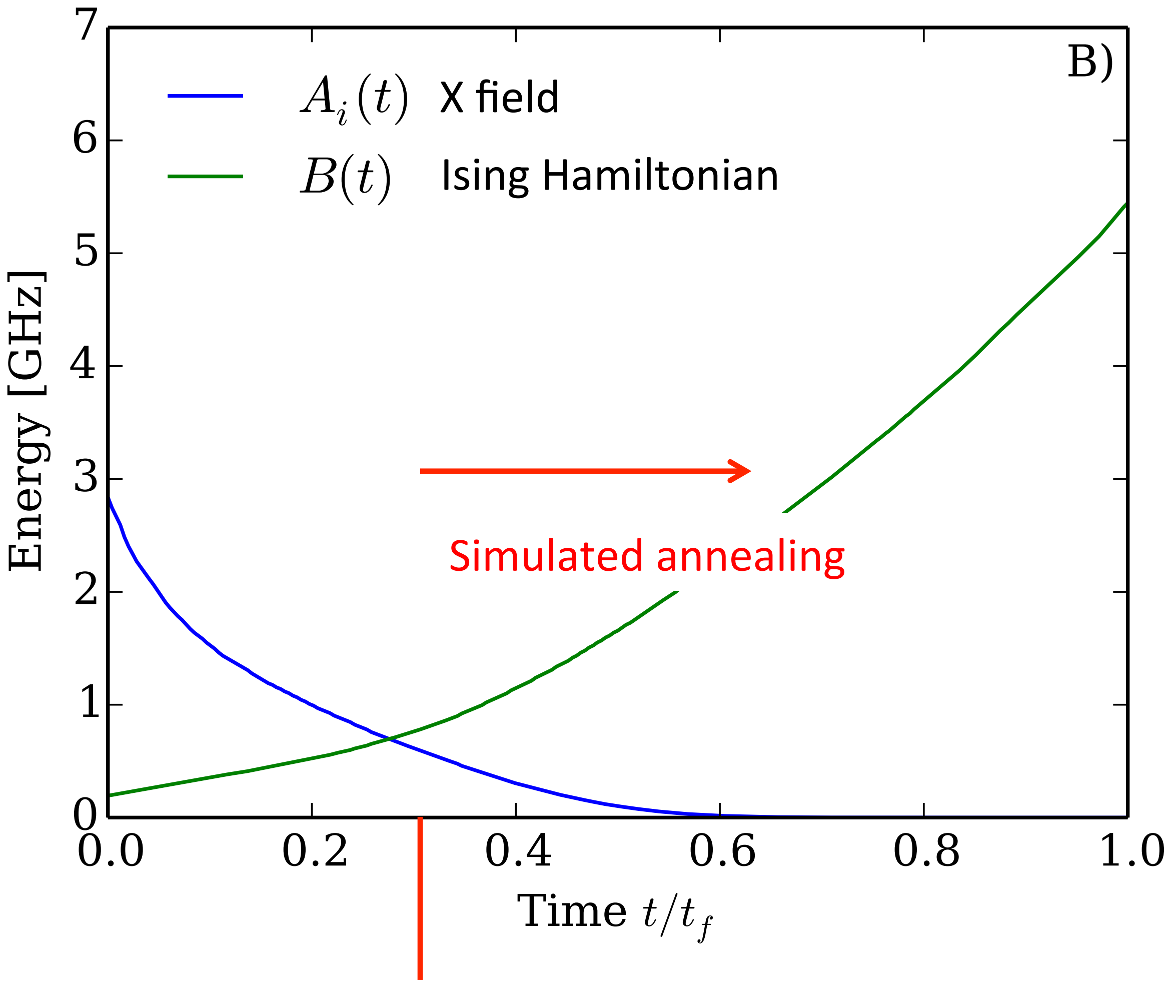

Furthermore…

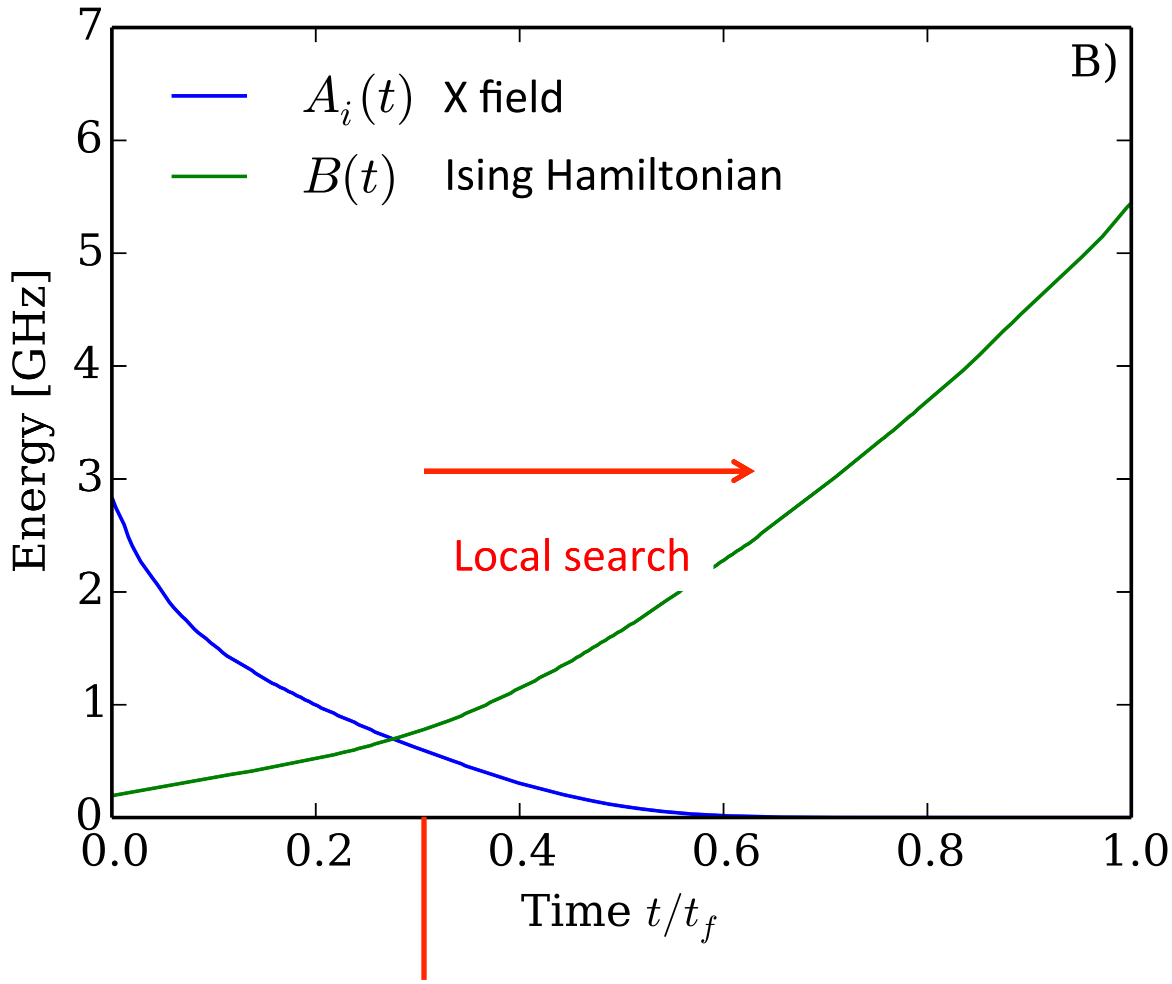

Furthermore…

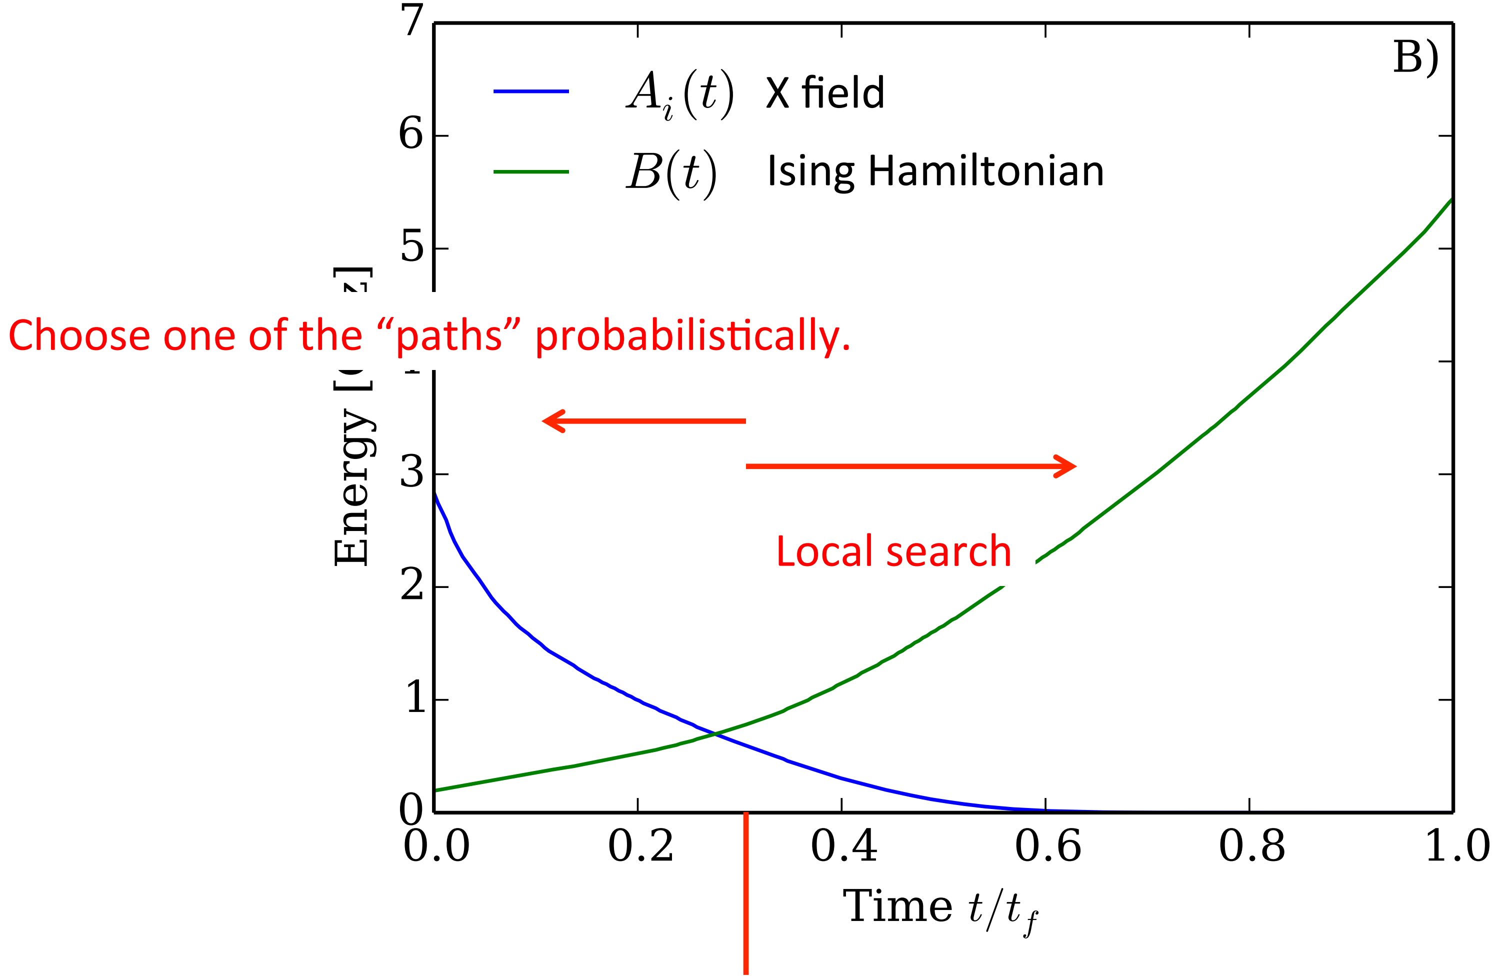

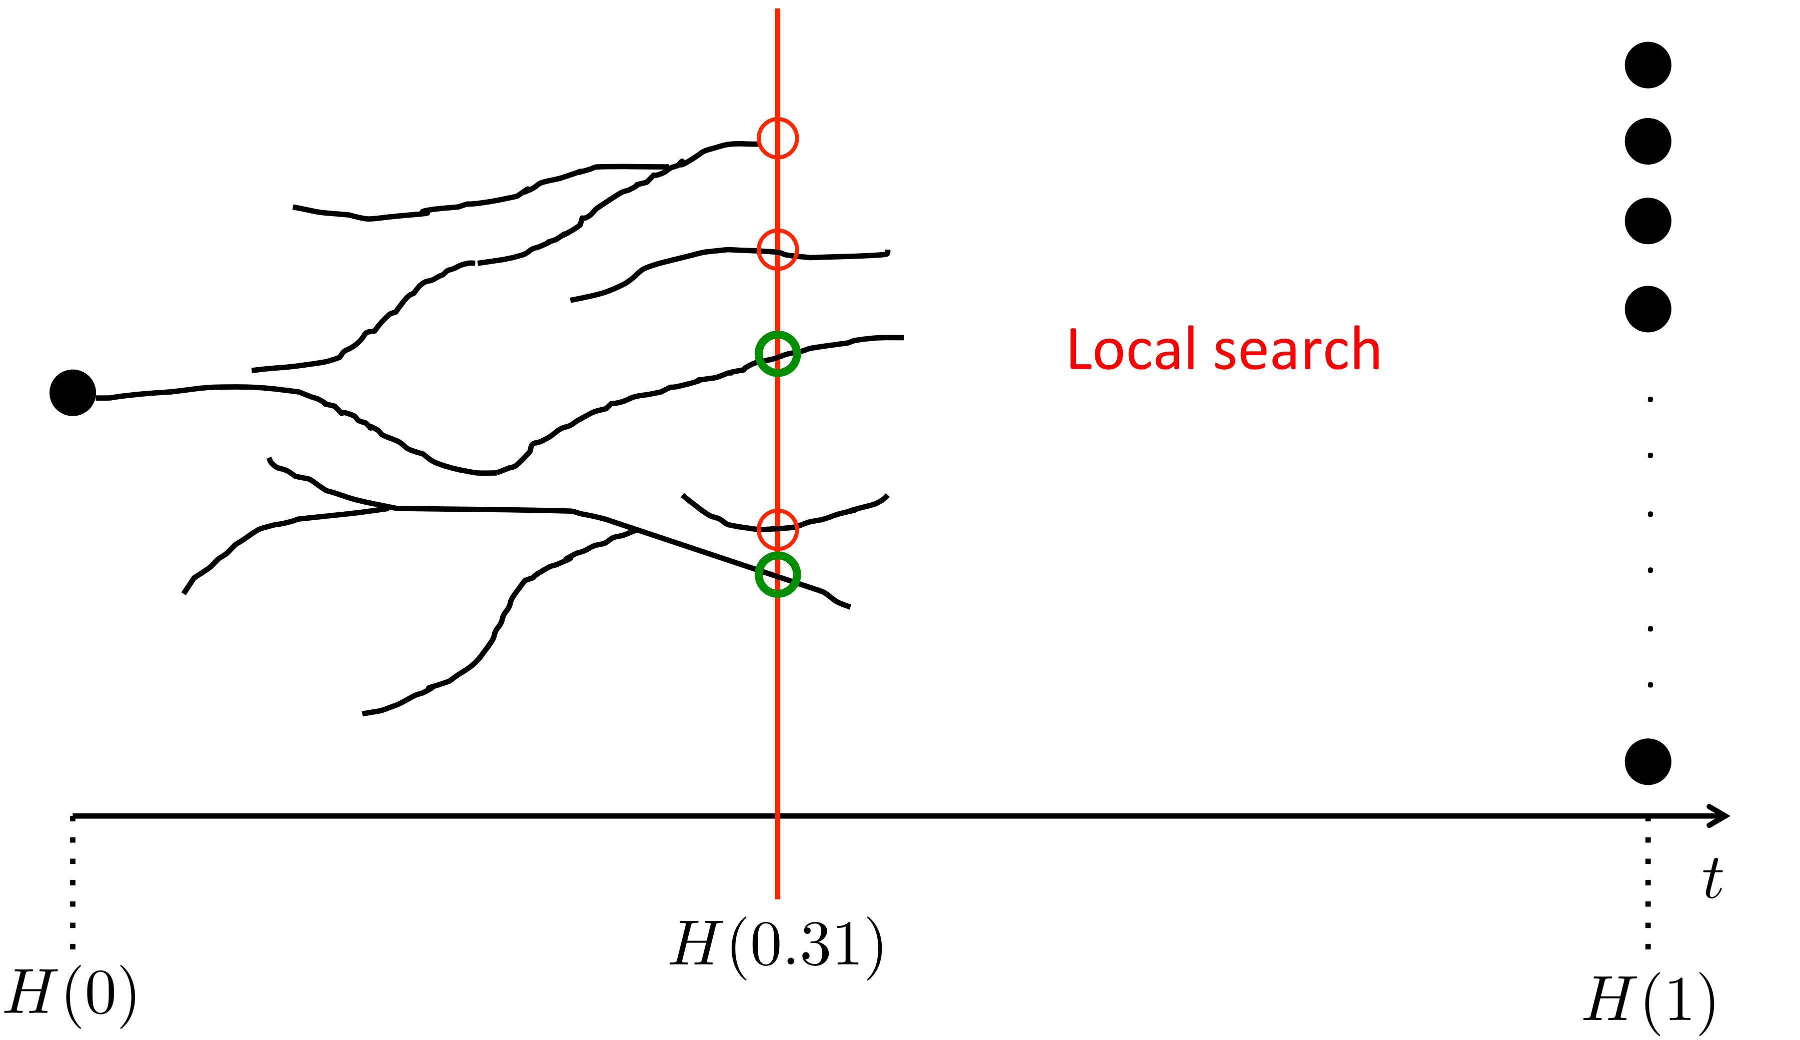

Local search
t
$H(0)$
$H(0.31)$
$H(1)$

How is it possible that exploring such a small search space, this model is still capable of solving a 108-bit problem so well?

There are 2^{108} possible solutions, but we are finding the right solution by just looking at a dozen of them!

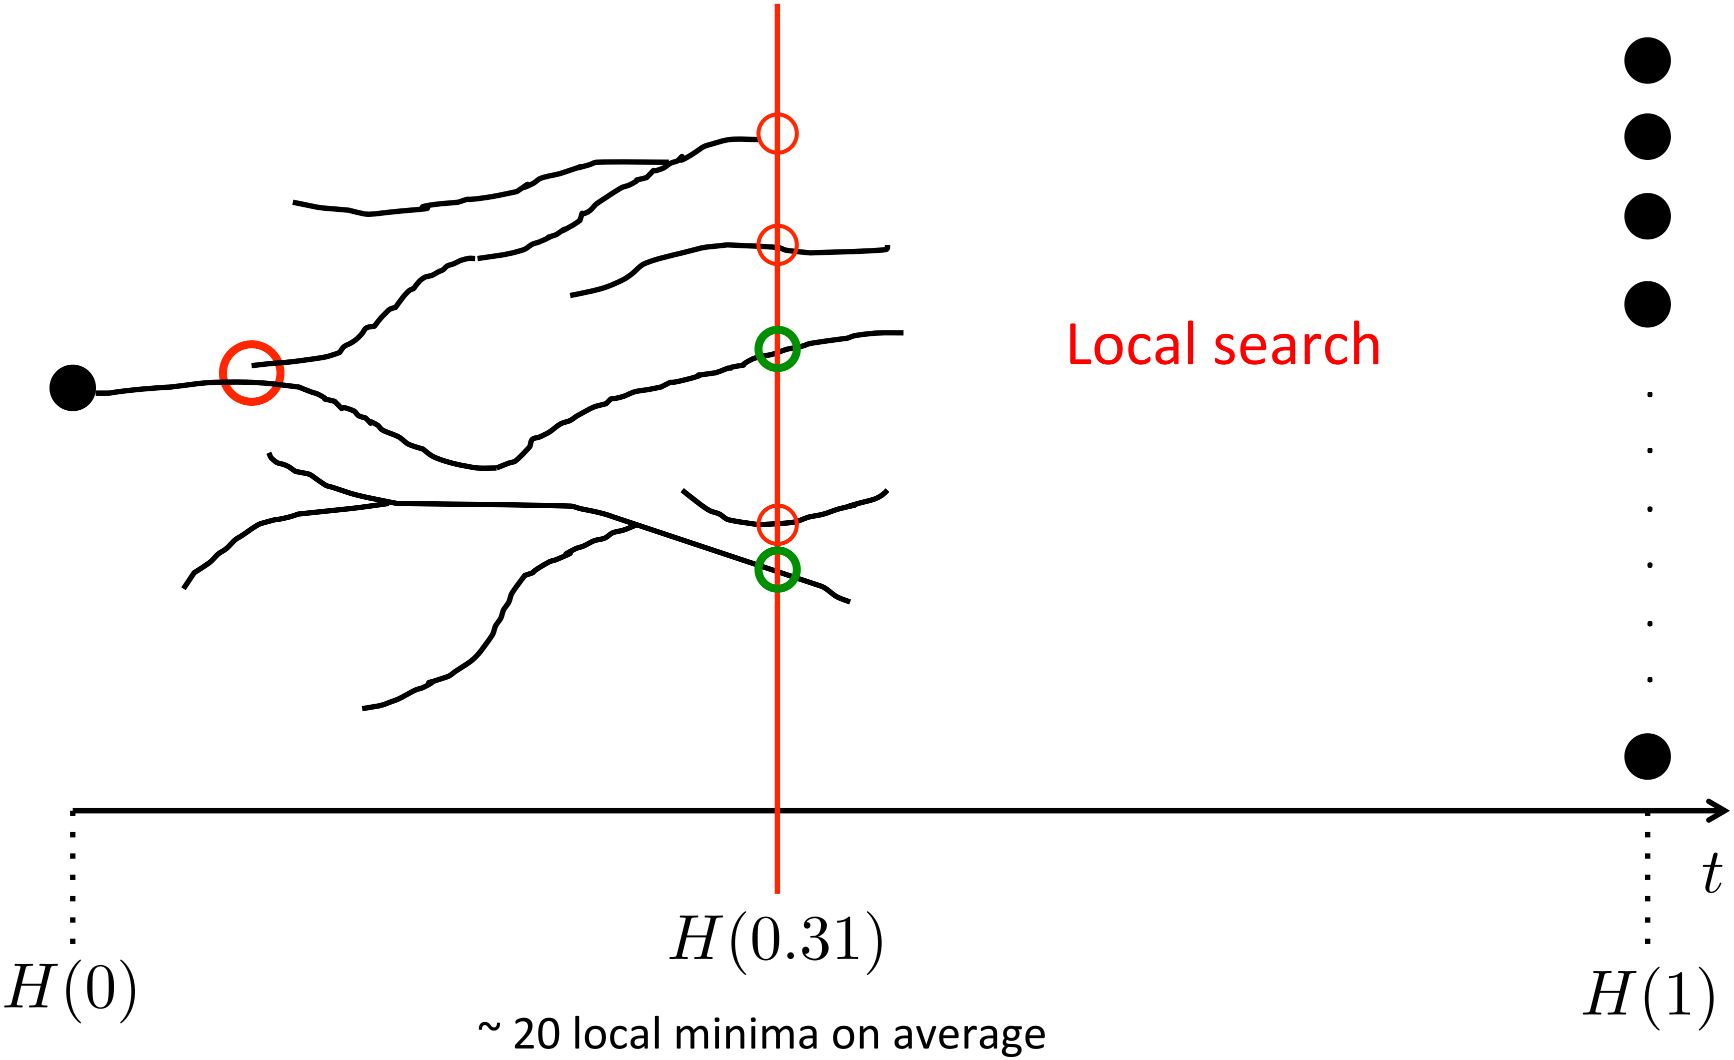

What choice is being made?

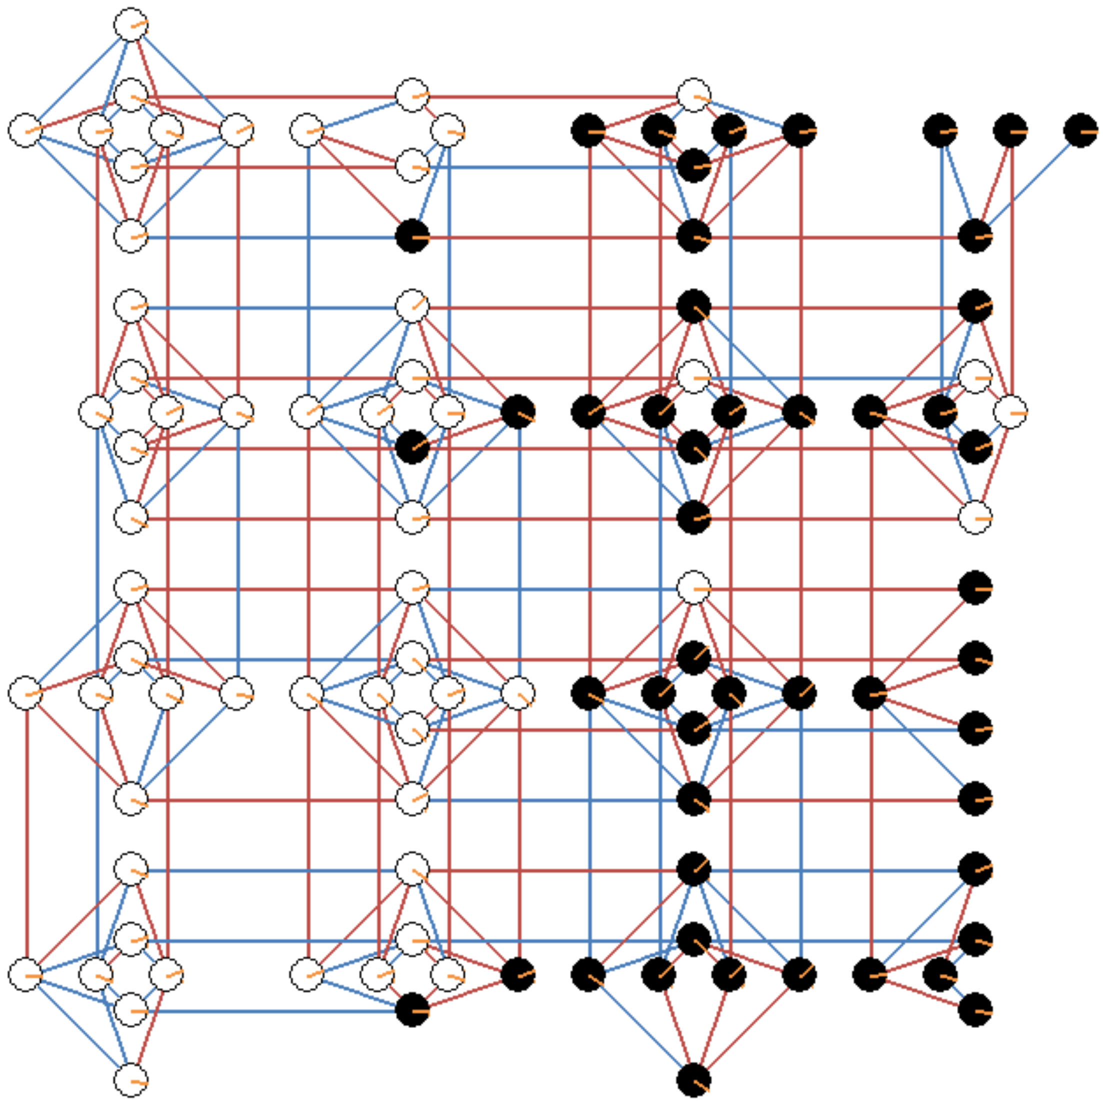

What choice is being made?

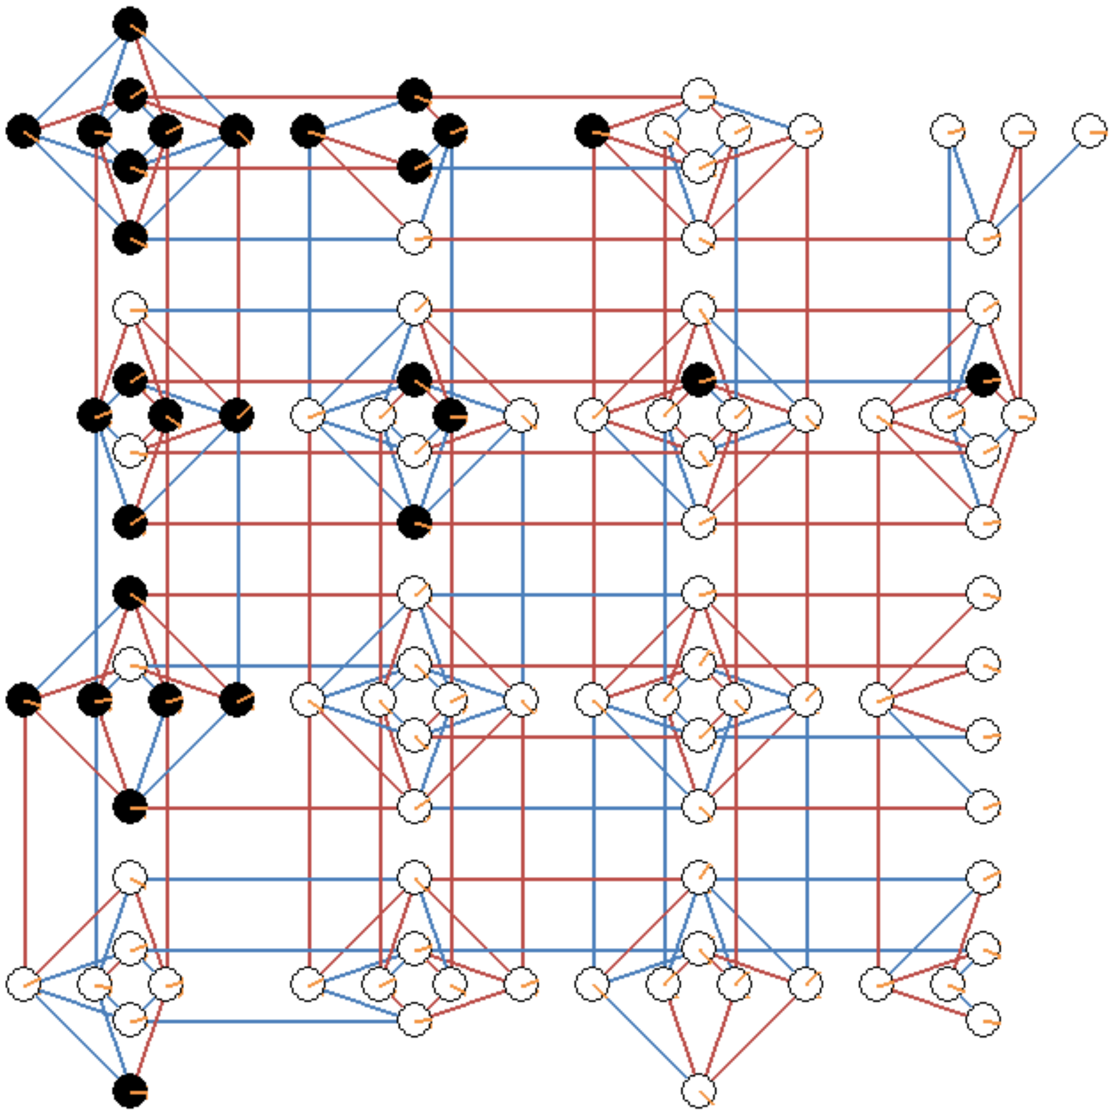

What choice is being made?

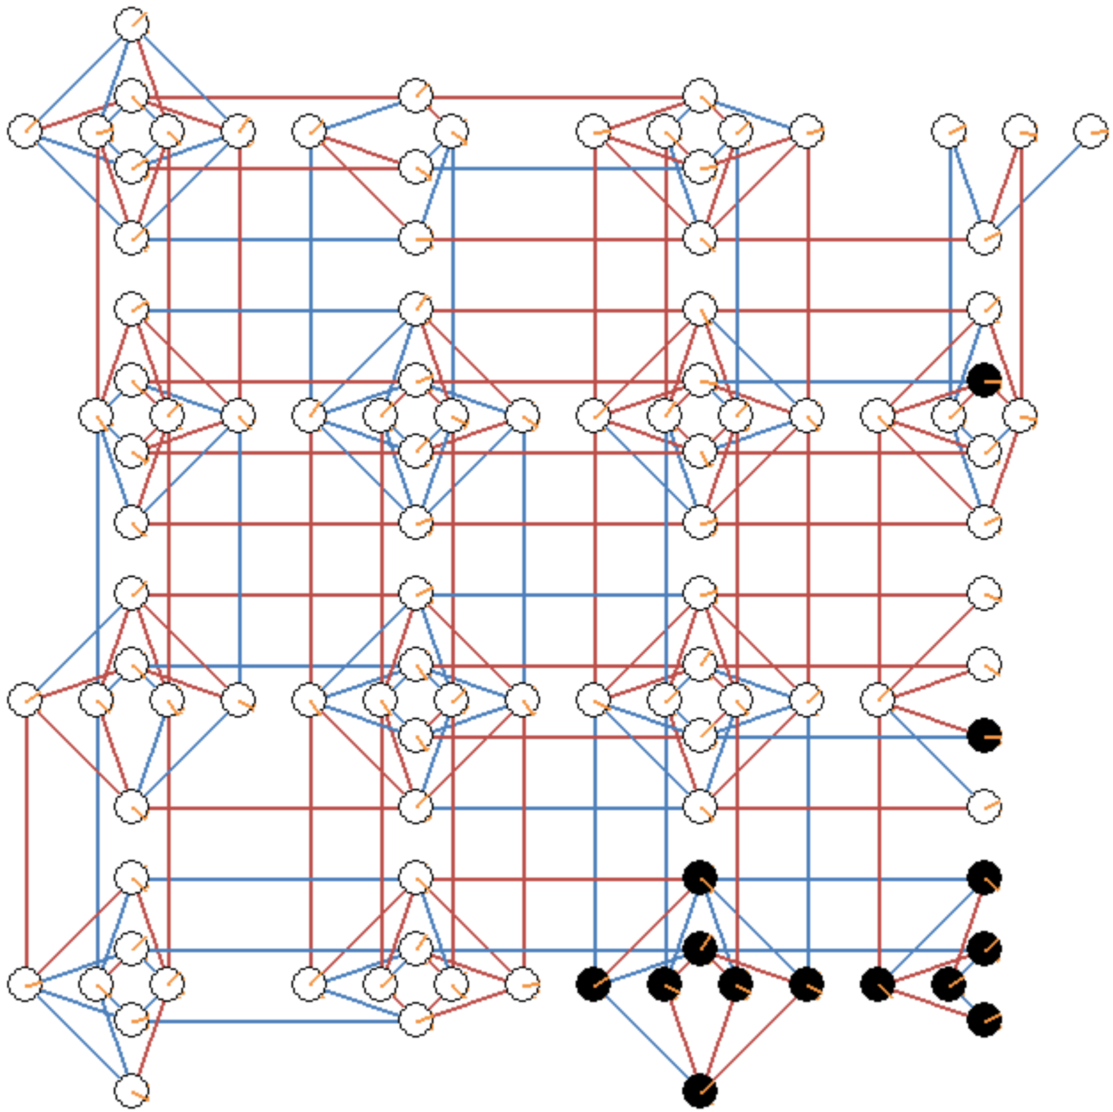

What choice is being made?

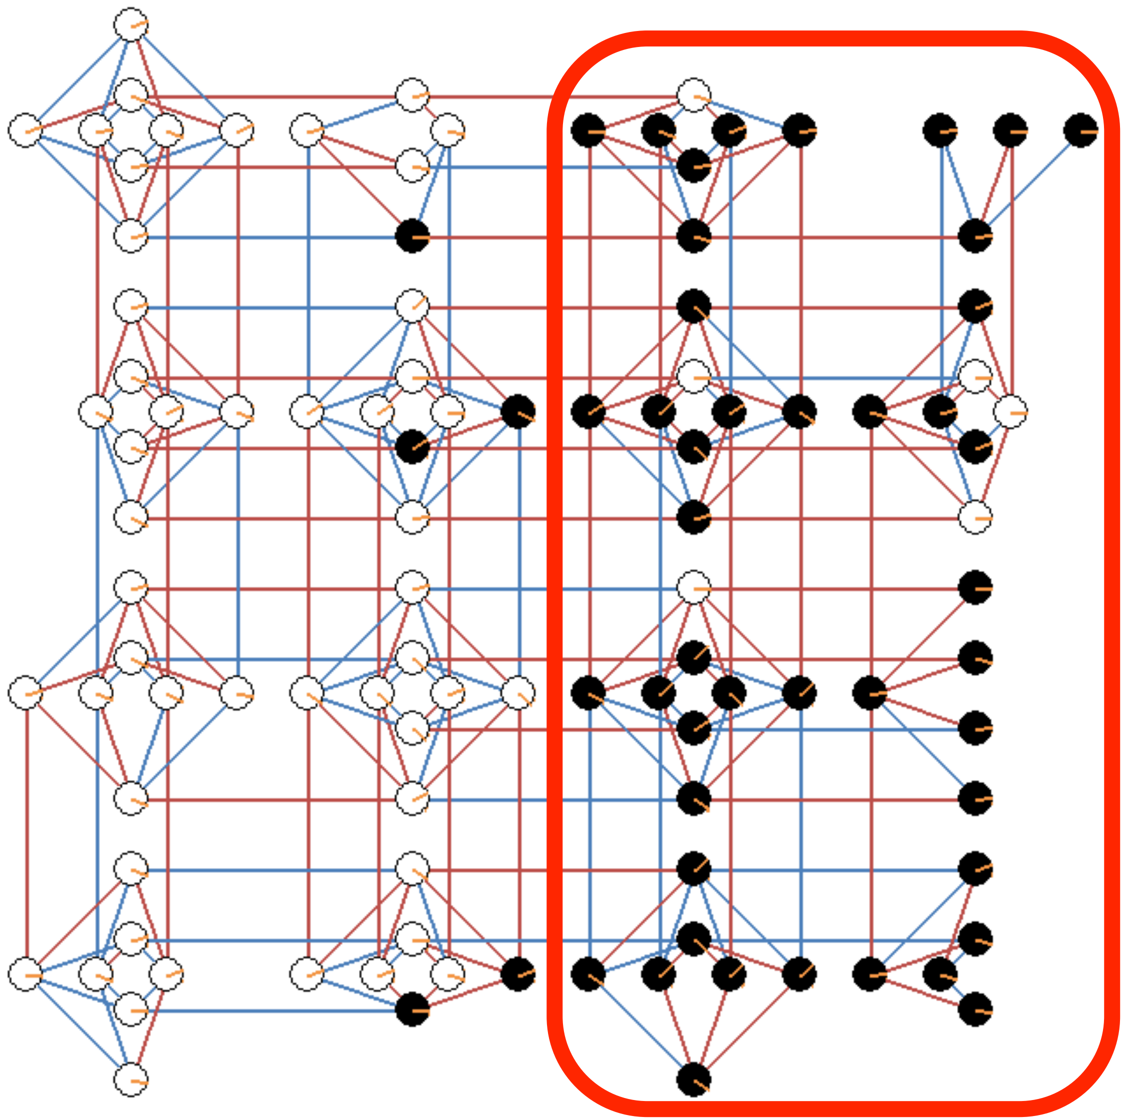

What choice is being made?

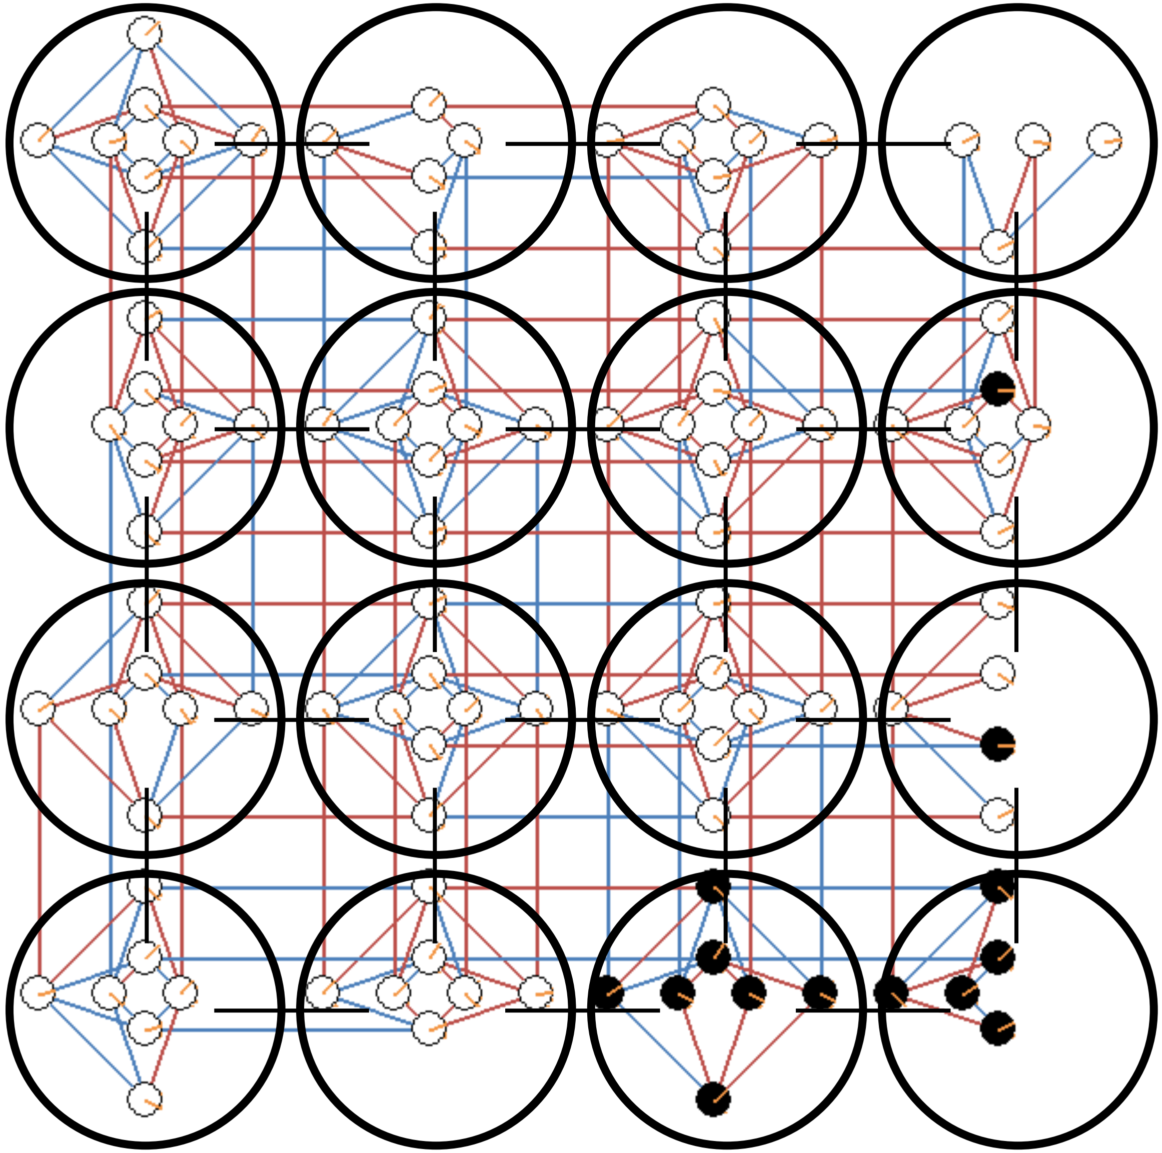

What choice is being made?

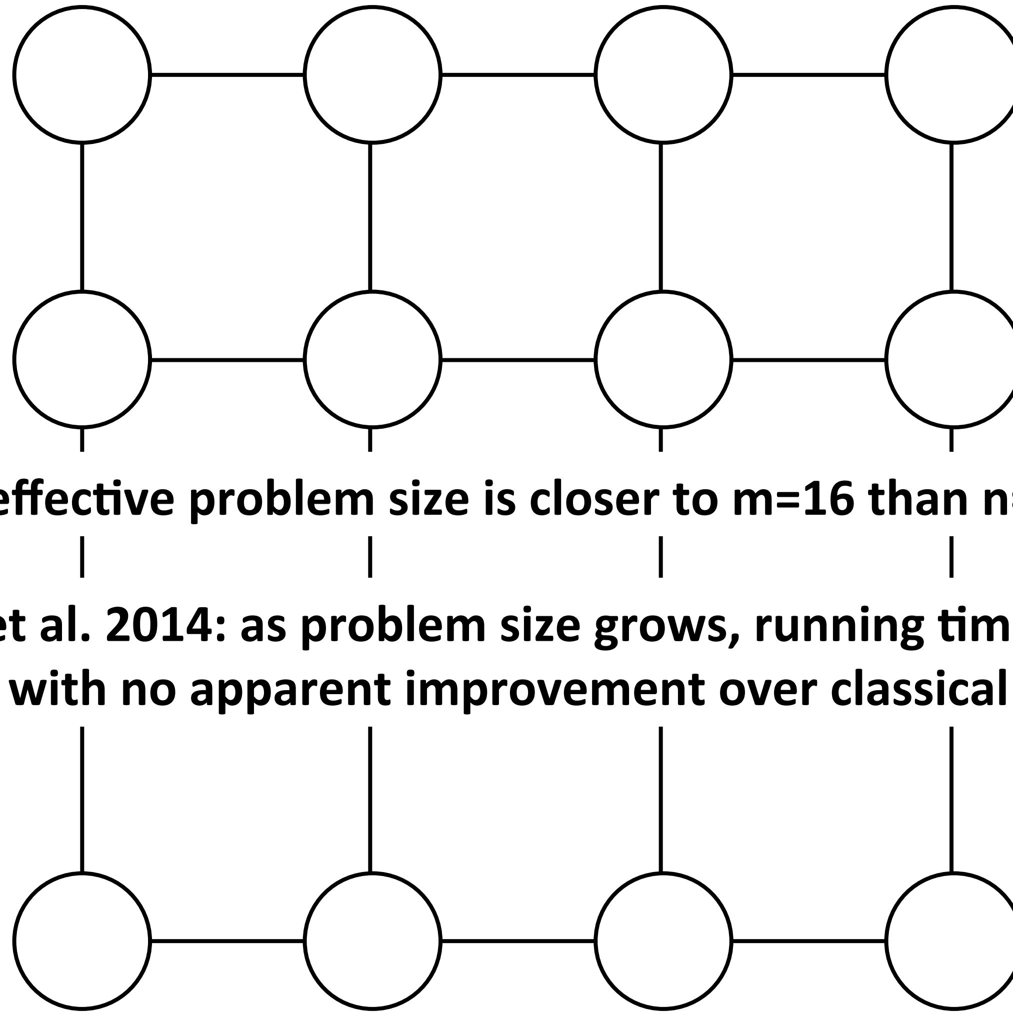

The effective problem size is closer to m=16 than n=108?

Rønnow et al. 2014: as problem size grows, running time scales exponentially with no apparent improvement over classical algorithms.

Quantum tunneling? vs. Classical thermal effects?

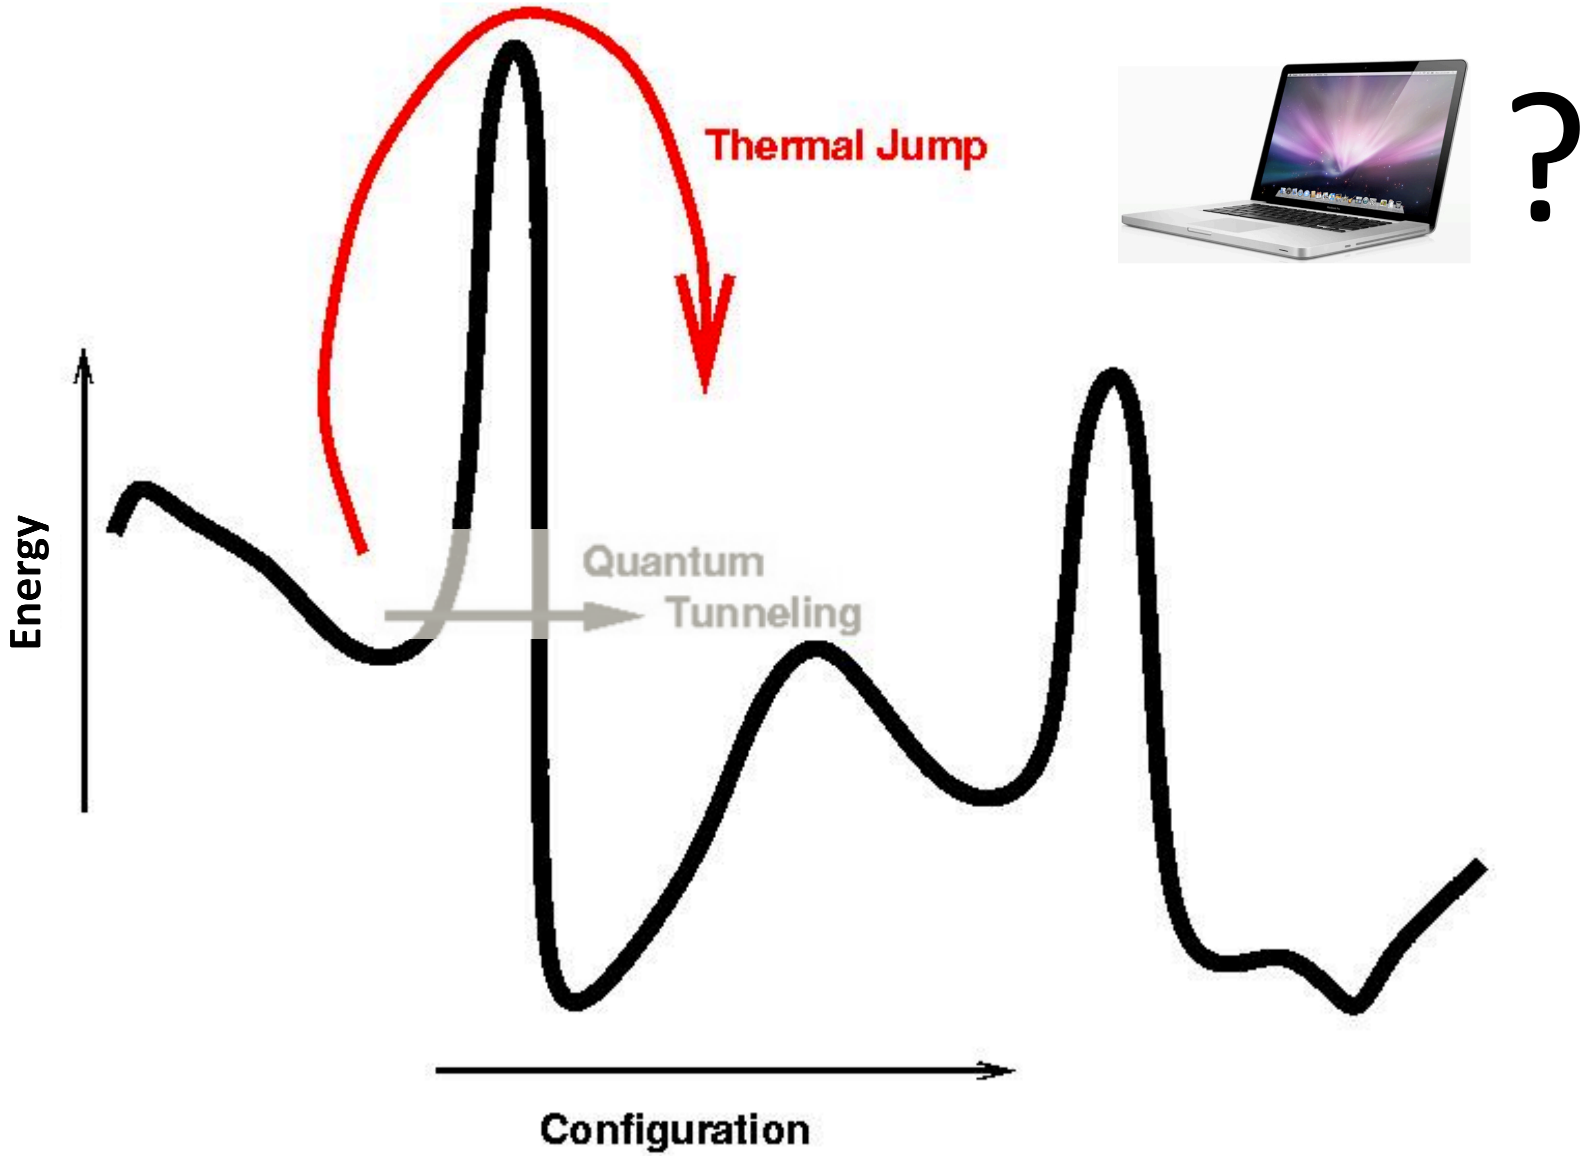

Image source: wiki/quantum_annealing

A tiny wee bit about recent developments…

- Vinci et al. (2014) recently uploaded a preprint (arXiv:1403.4228) in which they propose an experiment that distinguishes between our model and the D-Wave machine.
- The experiment involves local z-fields.

$$H(t) = -A(t)\sum_i \sin\theta_i - B(t)\left(\sum_{i\sim j} J_{ij}\cos\theta_i\cos\theta_j + \sum_i h_i\cos\theta_i\right)$$

- The discrepancy between the experiment and our model seems to stem from first-order vs. second-order terms in H.
- Preliminary investigations suggest that calibration of the local z-fields h_i plays an important role.
- When we add some noise to this calibration, our model seems to show similar behavior to the machine.
- ***However***, it is not clear whether this is an appropriate question for our "0th-order" model.

A tiny wee bit about recent developments…

- Vinci et al. (2014) recently uploaded a preprint (arXiv:1403.4228) in which they claim to have refuted our model using a different type of experiment.

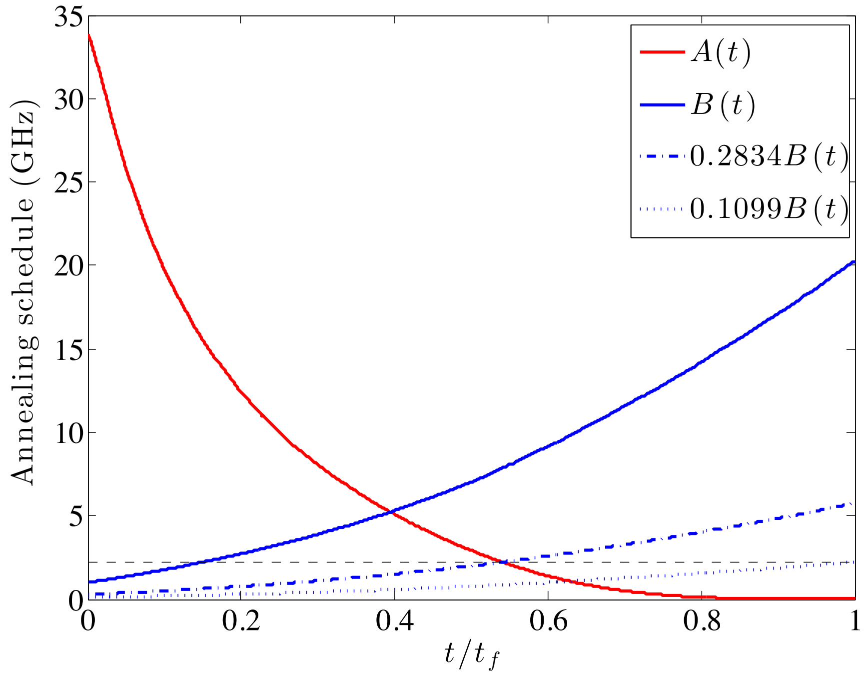

They report that our model does not reproduce the machine's behavior when the final Hamiltonian is very weak compared to the X-field.

But…

As the final Hamiltonian is turned down, the effect of noise becomes more and more significant, which means we are more into the classical regime. So in a sense, our model is behaving **too quantumly**…?

A tiny wee bit about recent developments…

- Vinci et al. (2014) recently uploaded a preprint (arXiv:1403.4228) in which they claim to have refuted our model using a different type of experiment.
- Nonetheless, our preliminary investigations seem to suggest that under a reasonable assumption on calibration errors, our model *does reproduce* the machine's behavior.
- **Yet…** is it really fair to ask such questions to our model?

A "0th-order" model

(In the sense that we made no attempt to model anything beyond the computational concept.)

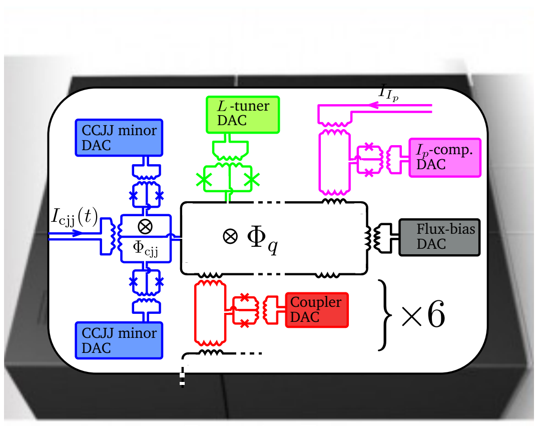

Cannot expect this simple model to explain everything.

If one wants to explain everything, should model the inside of the box.

Key question should be, "Can one demonstrate some *computationally meaningful* quantum phenomenon that our model does not describe?"

Is D-Wave "quantum"?

"quantumness"

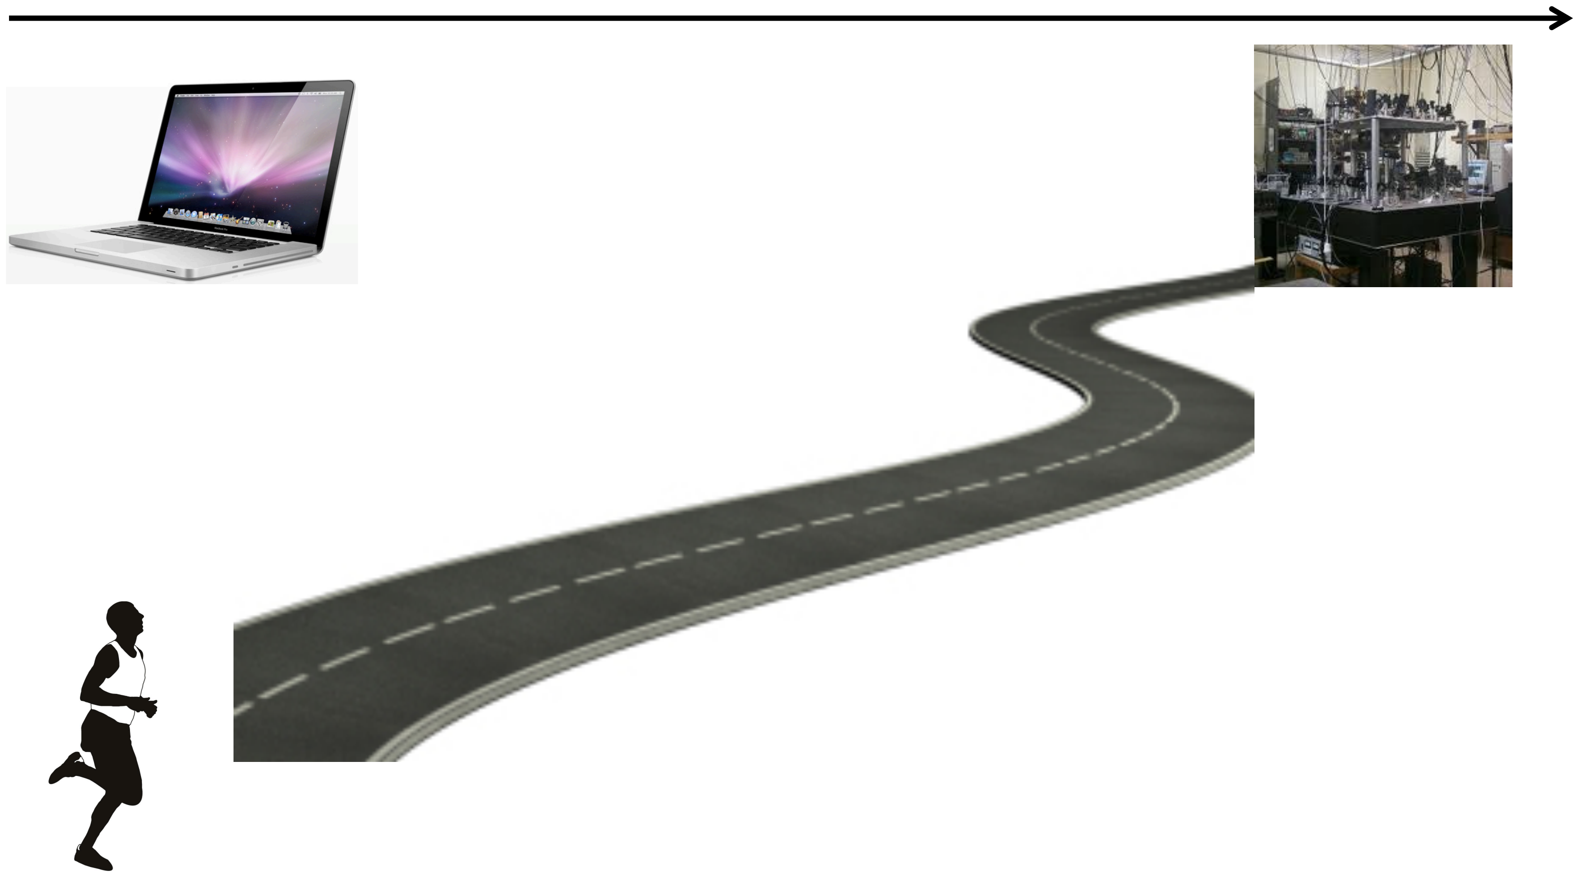

It is only a beginning . . .