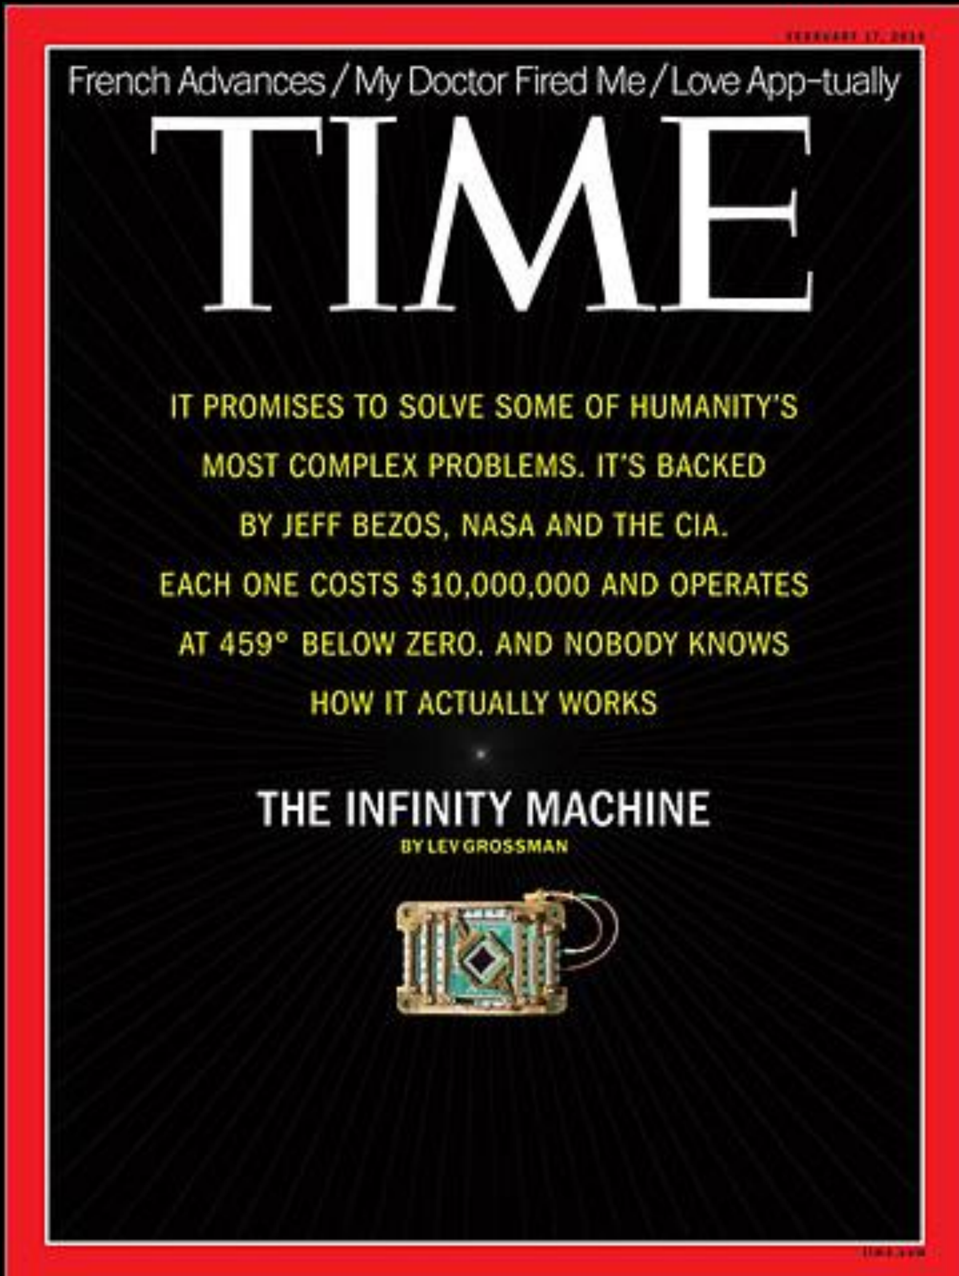

Adiabatic Quantum Optimization & D-Wave

Do we know how it works?
Does it really work?

[Time 2/2014, photo by Gregg Segal]

Adiabatic Quantum Optimization & D-Wave

Do we know how it works?
Does it really work?

TECH QUANTUM

The Quantum Quest for a Revolutionary Computer

Lev Grossman @leverus | Feb. 6, 2014

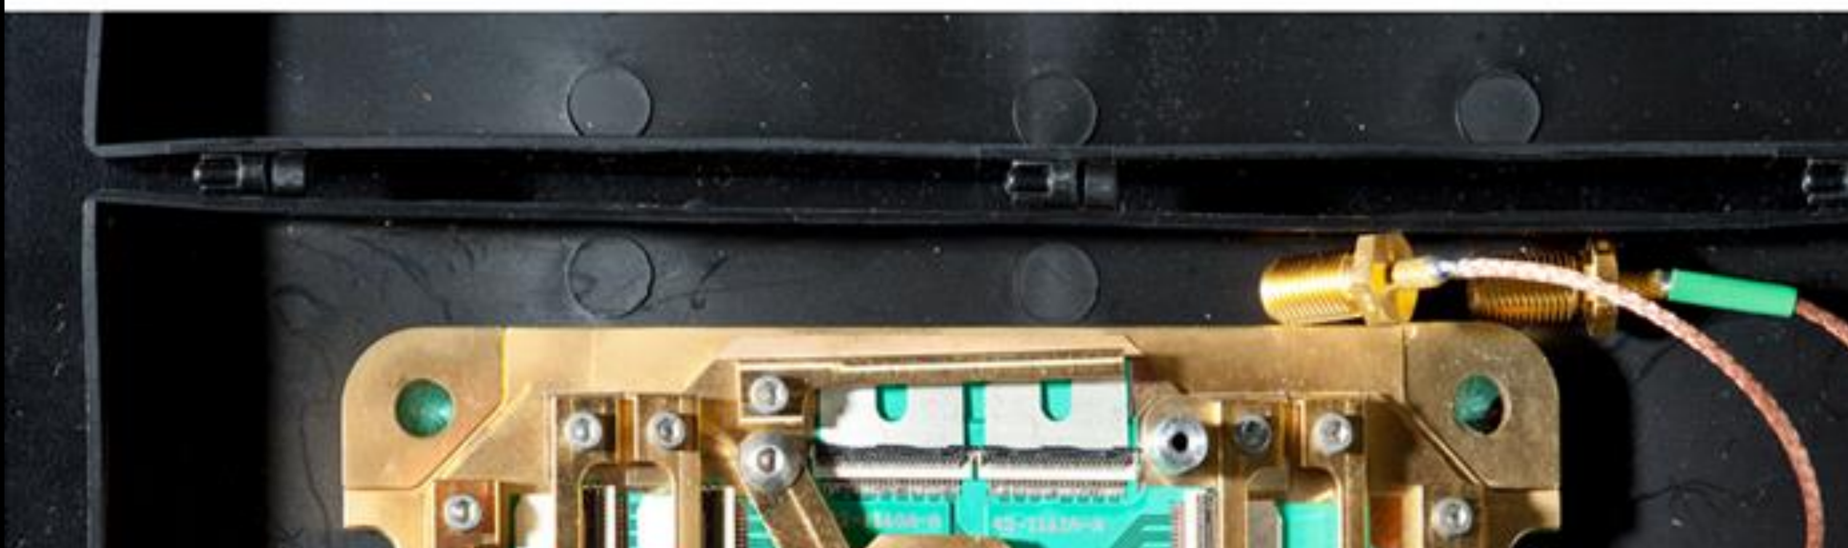

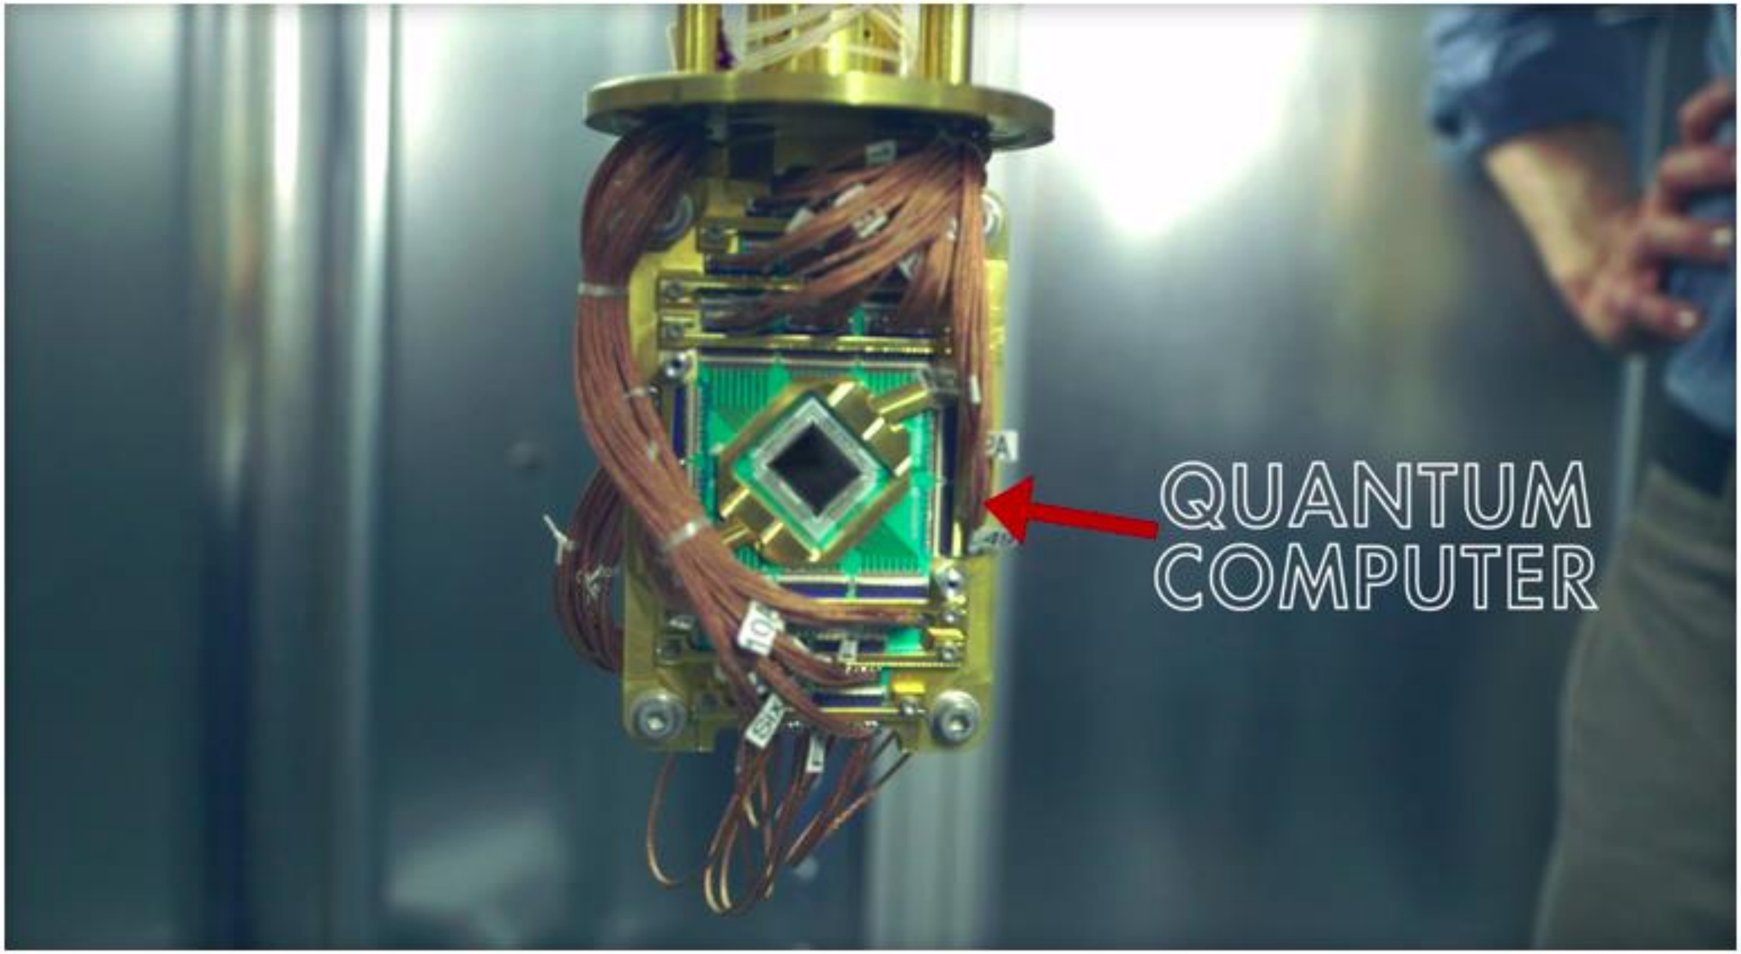

Screenshot by Nick Statt/CNET

503 qubits!

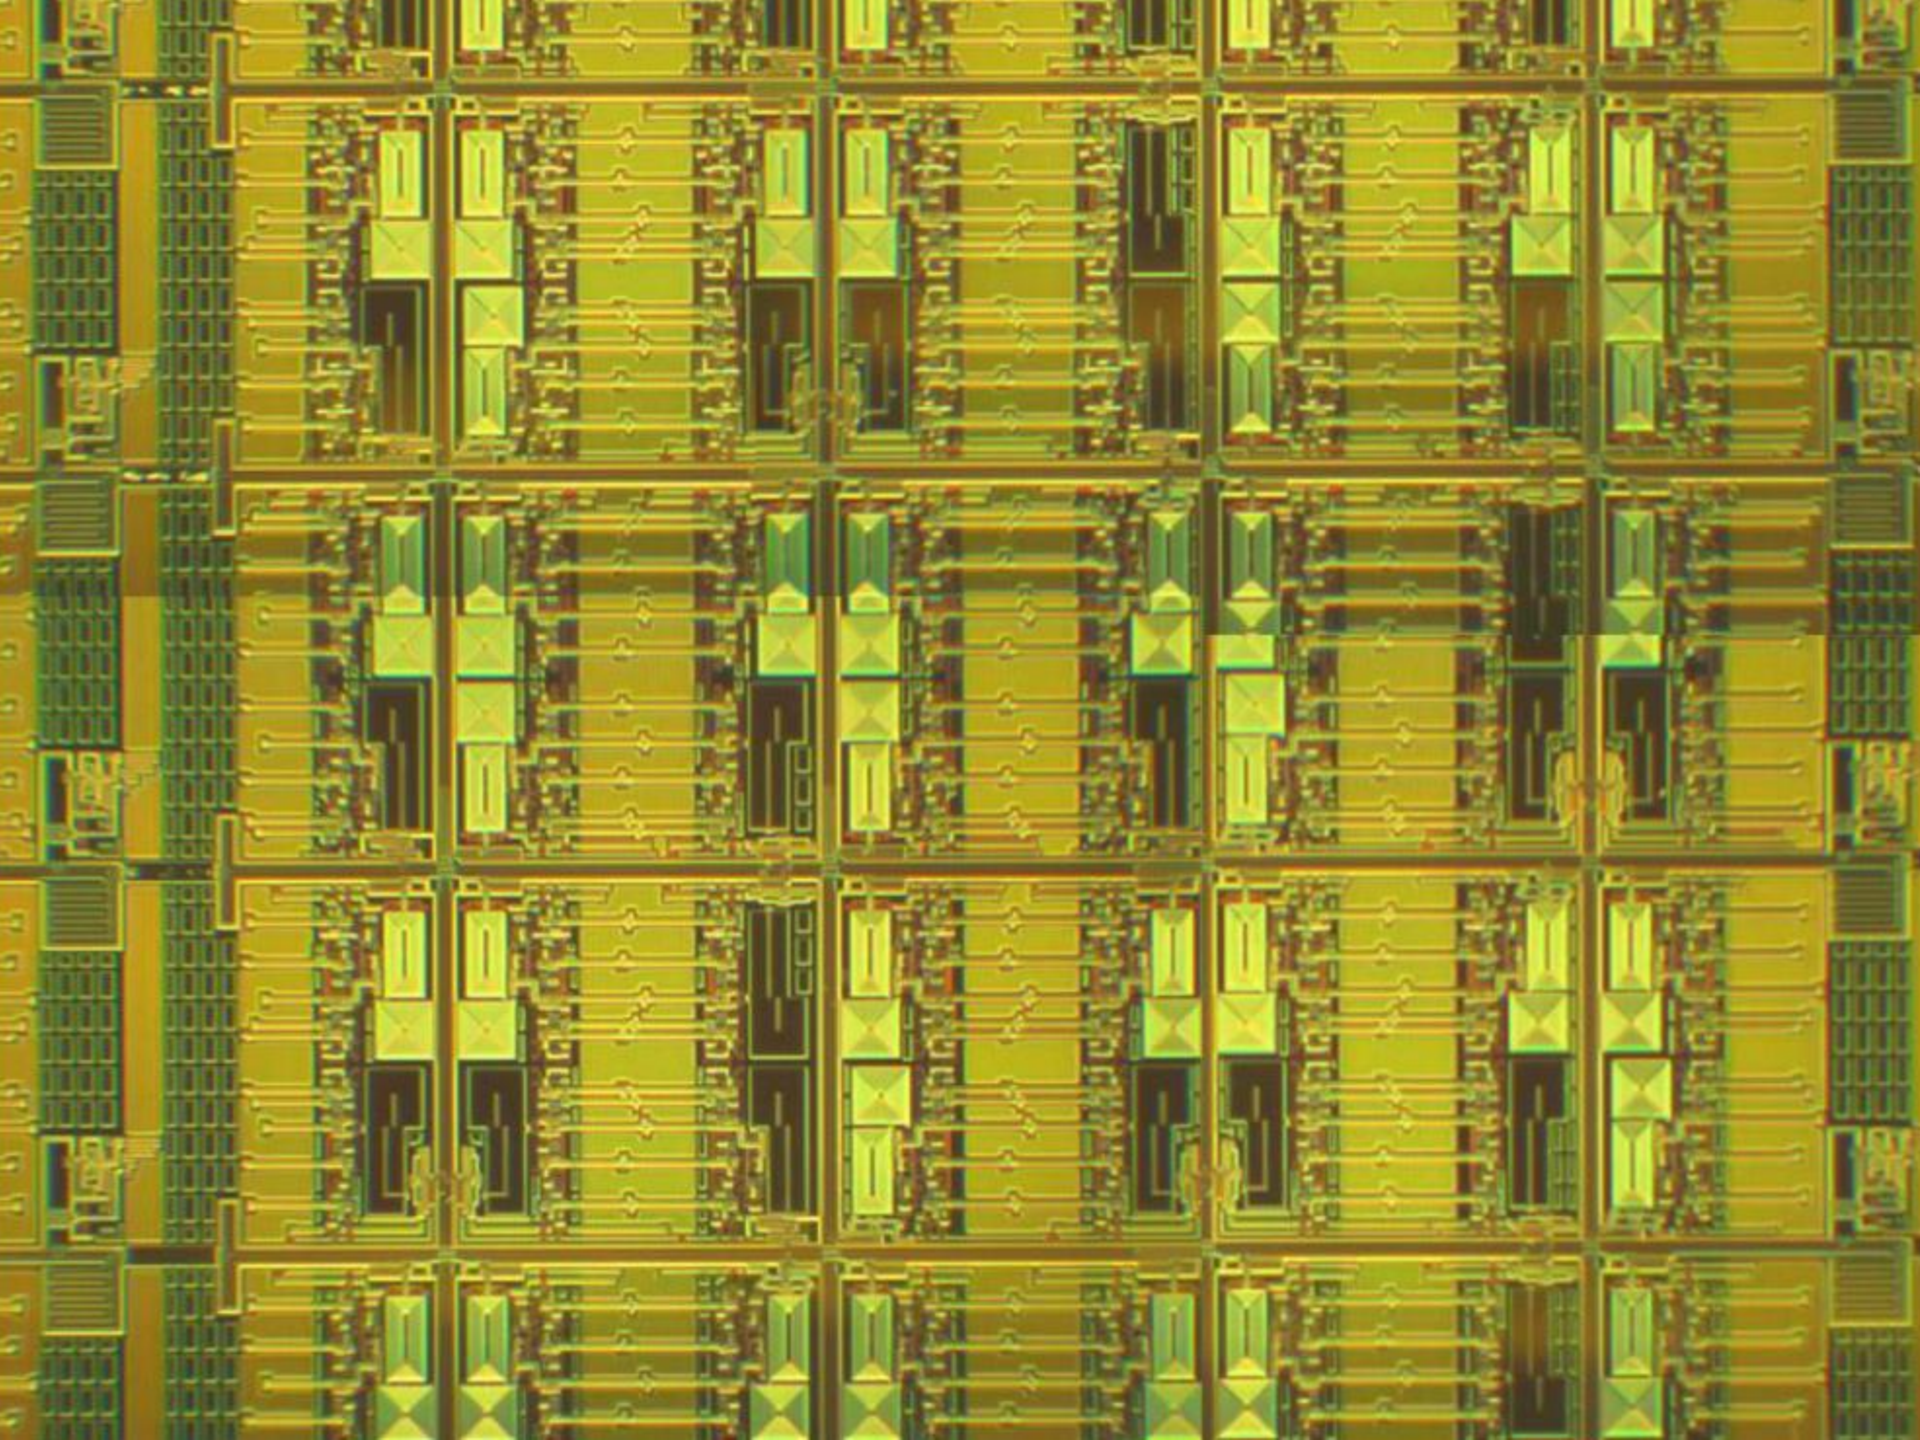

Commercial quantum computer leaves PC in the dust

› 12:18 10 May 2013 by Jacob Aron
› For similar stories, visit the Quantum World Topic Guide

For the first time, a commercially available quantum computer has been pitted against an ordinary PC – and the quantum device left the regular machine in the dust.

D-Wave, a company based in Burnaby, Canada, has been selling quantum computers since 2011, although critics expressed doubt that their chips were actually harnessing the spooky action of quantum mechanics. That's because they use a non-mainstream method called adiabatic quantum computing.

Unlike classical bits, quantum bits, or qubits, can take the values 0 and 1 at the same time, theoretically offering much faster computing speed. To be truly quantum, the qubits must be linked via the quantum property of entanglement. That's impossible to measure while the device is operating. But in March, two separate tests of the D-Wave device showed indirect evidence for entanglement.

Now Catherine McGeoch of Amherst College, Massachusetts, a consultant to D-Wave, has put their computer through its paces and shown that it can beat regular machines. The D-Wave hardware is designed to solve a particular kind of optimisation problem: minimising the solution of a complicated equation by

Commercial quantum computer leaves PC in the dust

McGeoch and her colleague Cong Wang of Simon Fraser University, in Burnaby, ran the problem on a D-Wave Two computer, which has 439 qubits formed from superconducting niobium loops. They also tried to solve the problem using three leading algorithms running on a high-end desktop computer. The D-Wave machine turned out to be around 3600 times faster than the best conventional algorithm.

is it useful
quantum

Hamiltonians

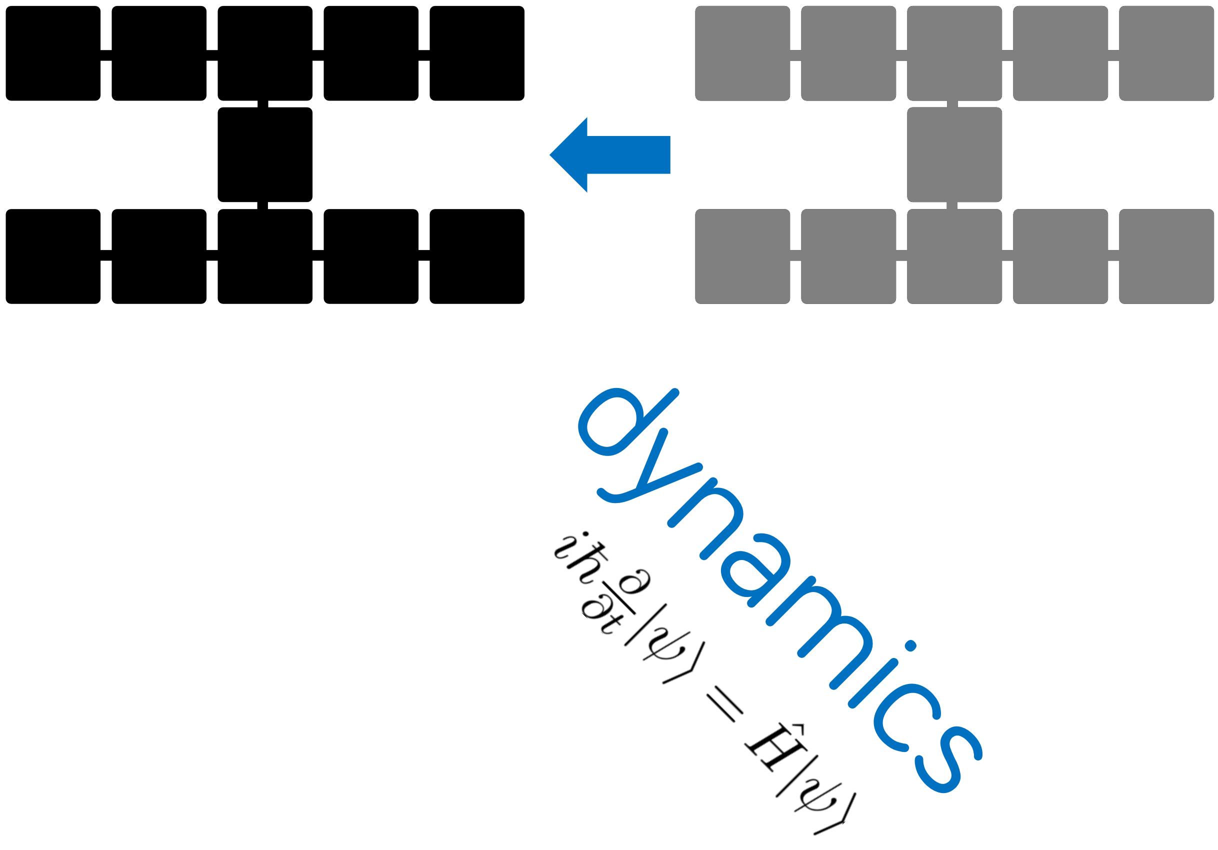
dynamics
$i\hbar\frac{\partial}{\partial t}|\psi\rangle = \hat{H}|\psi\rangle$

scattering

[fermilab]

scattering

Universal computation by multi-particle quantum walk

- dual-rail encoding
  N wavepackets

$$a_j^\dagger a_k + a_k^\dagger a_j$$

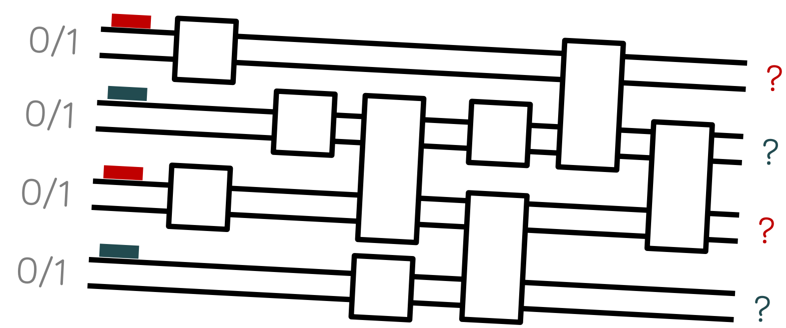

- CPHASE: interaction

$$a_j^\dagger a_k^\dagger a_j a_k$$

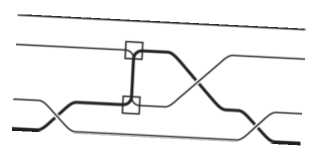

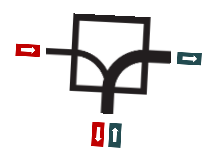

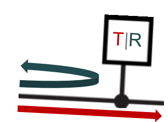

[Childs, Gosset, Webb, Science 339, 791 (2013)]

optimization

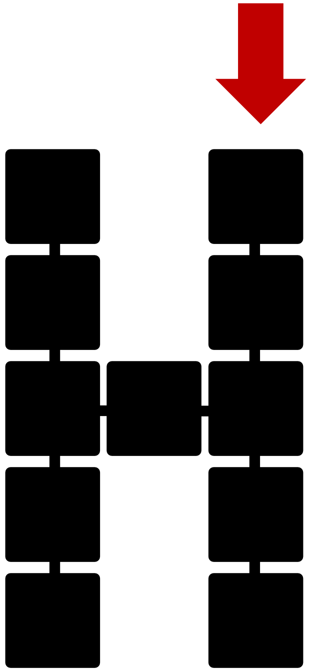

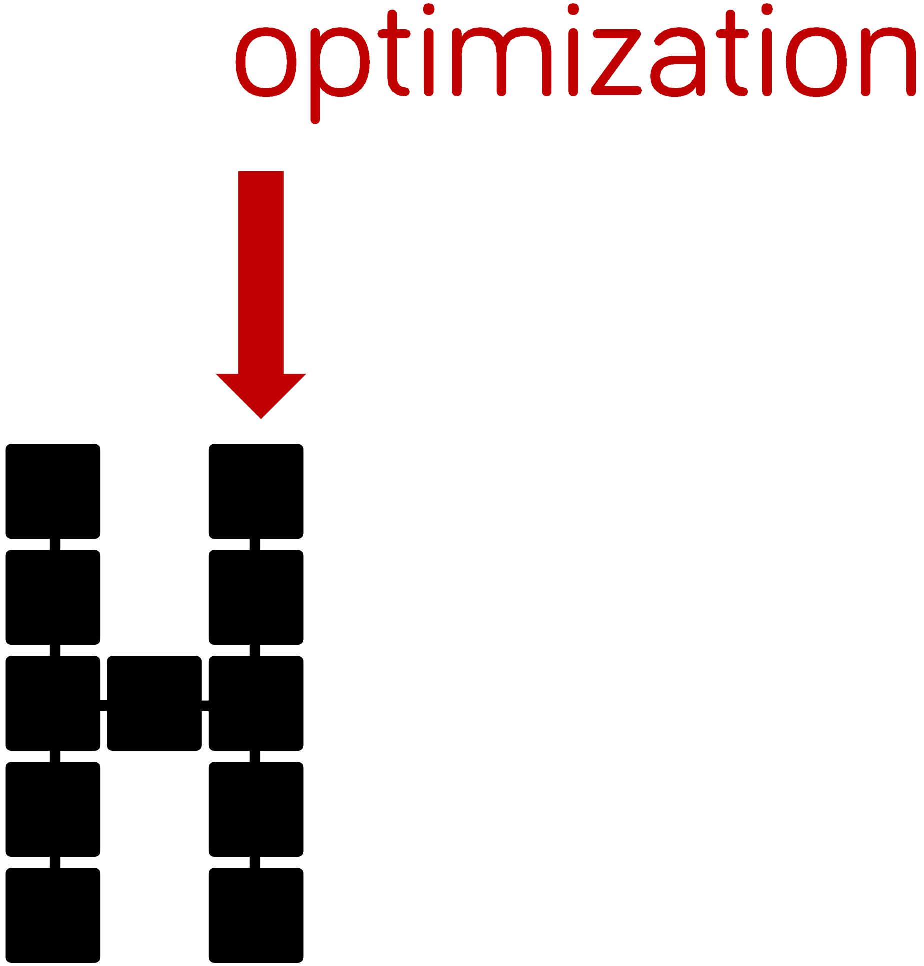
optimization

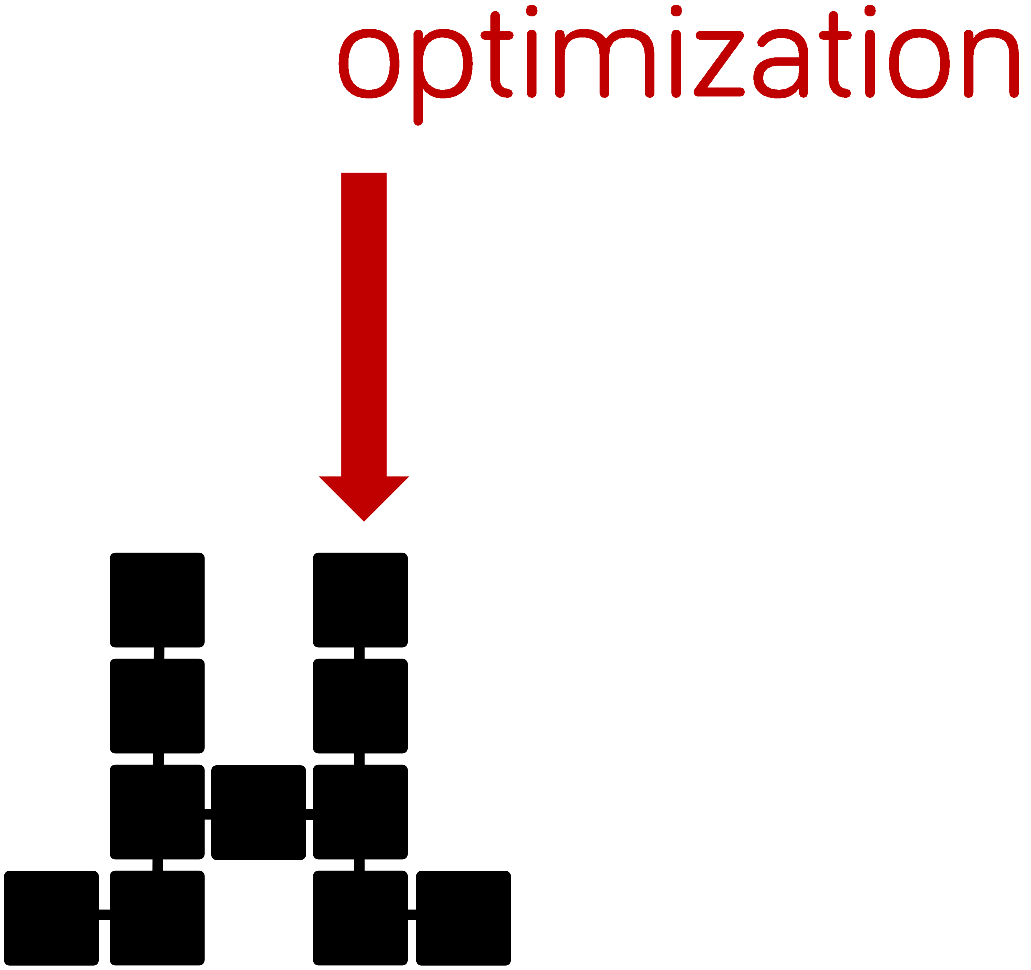
optimization

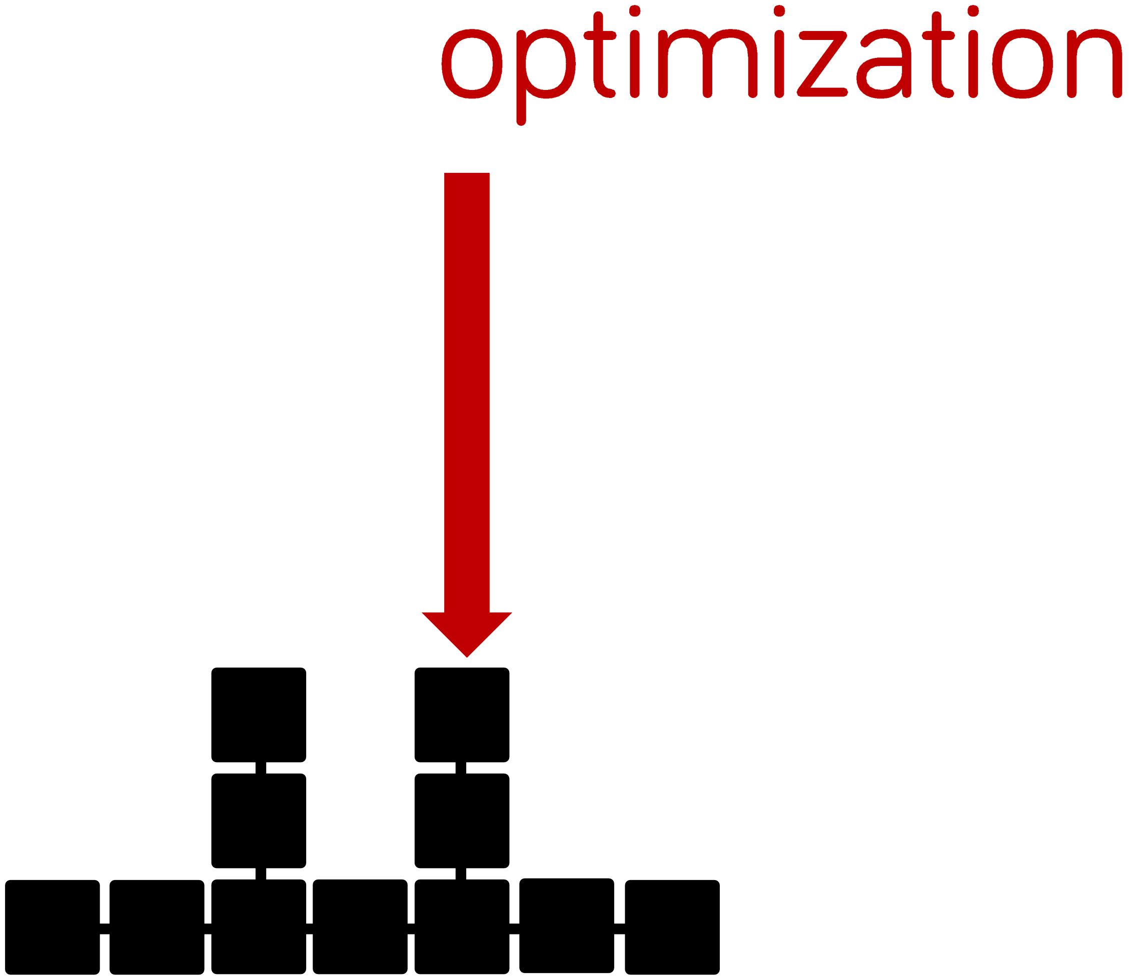
optimization

optimization

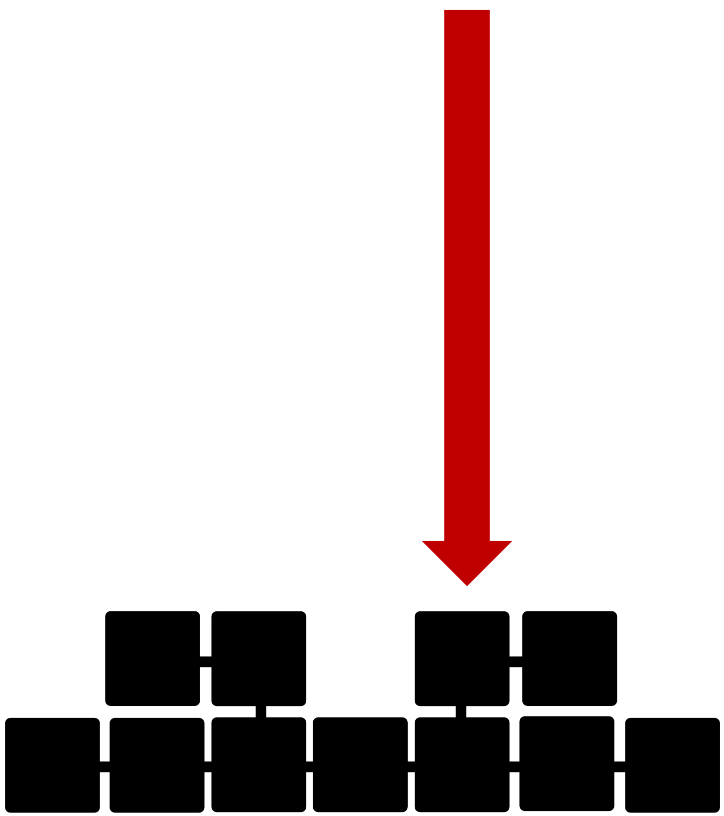

protein folding

spin glasses

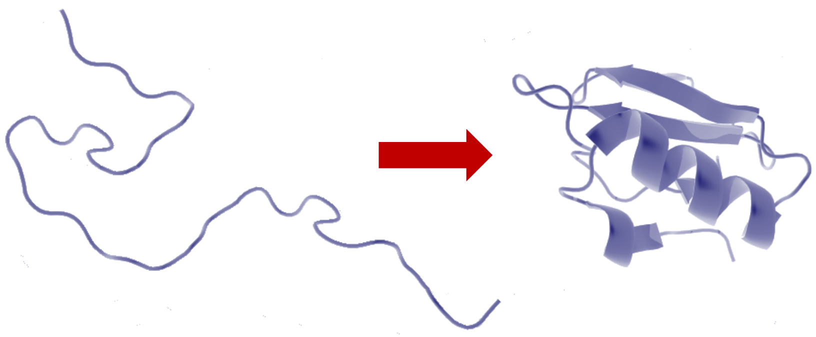

[wikipedia]

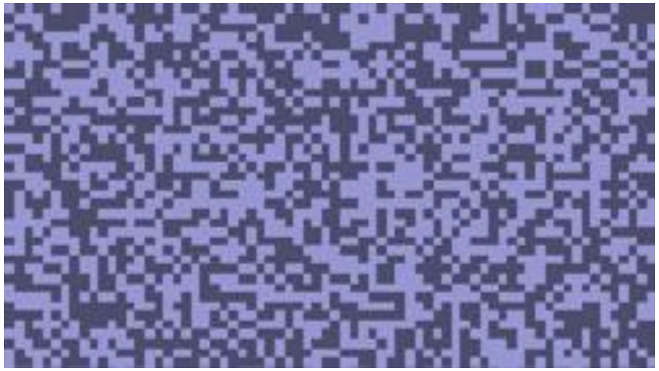

[uni-koeln.de]

local Hamiltonians

0 Frustrated systems

You can't make everybody happy.

antiferromagnetic
spin glass

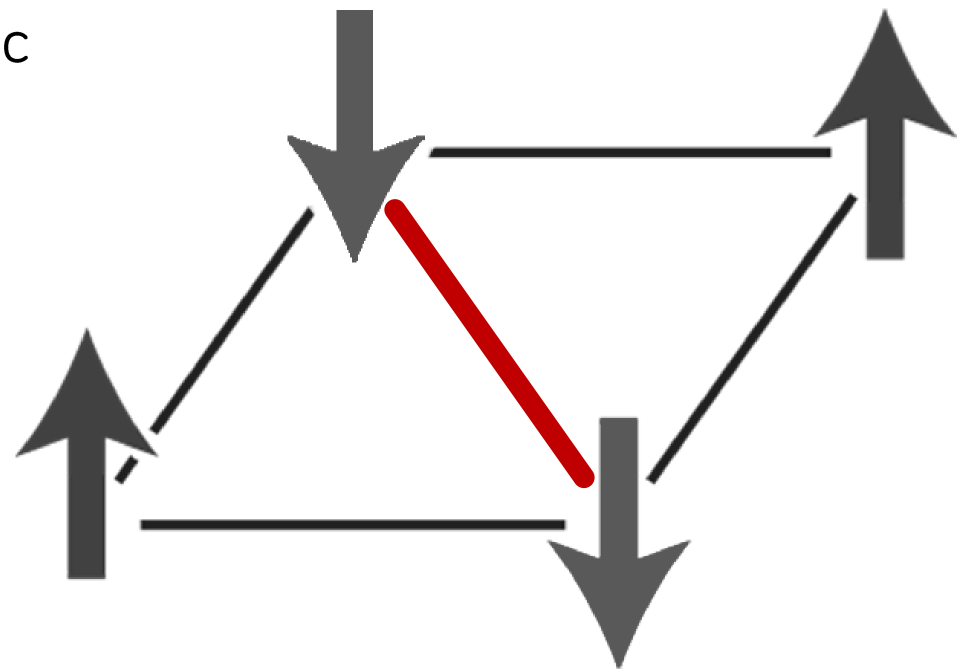

a global
ground state

HARD?

find & describe it?
is it entangled?


frustrated
FRUST
RATED

0 Local Hamiltonians

Many local questions.

- What are they like?

  ground state (energy)

  QMA-complete problems

$$H(t) = \sum_j H_j(t)$$

- What are they good for?

  quantum computing, chemistry, control, transport

  BQP universality

- local particle dimension 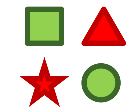
- interaction geometry 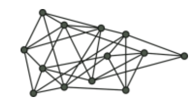
- time independence 
- translational invariance 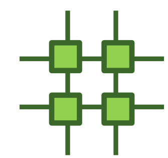
- promise and eigenvalue gaps 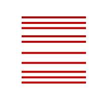
- energy × time cost 

how hard
is this
question

[J. Howard Miller]

P

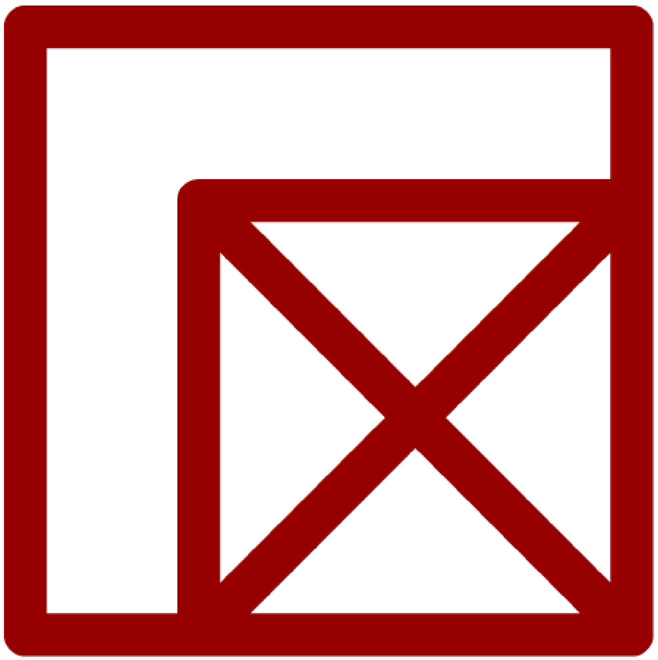

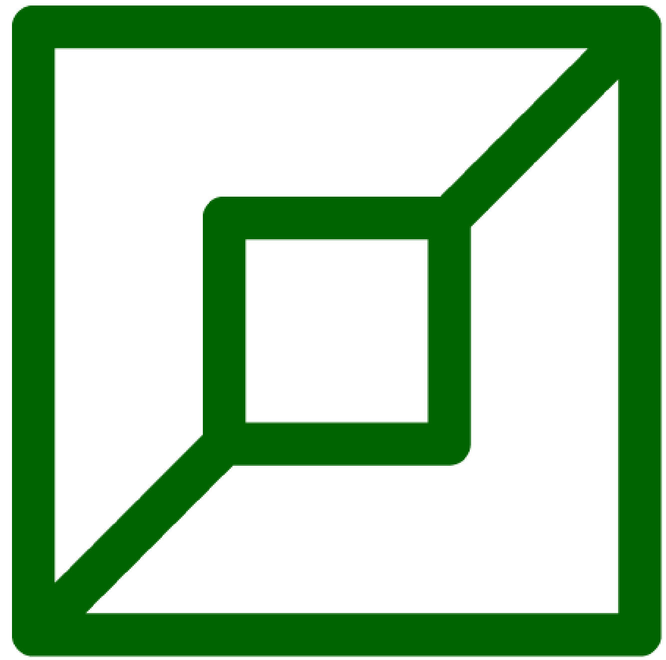

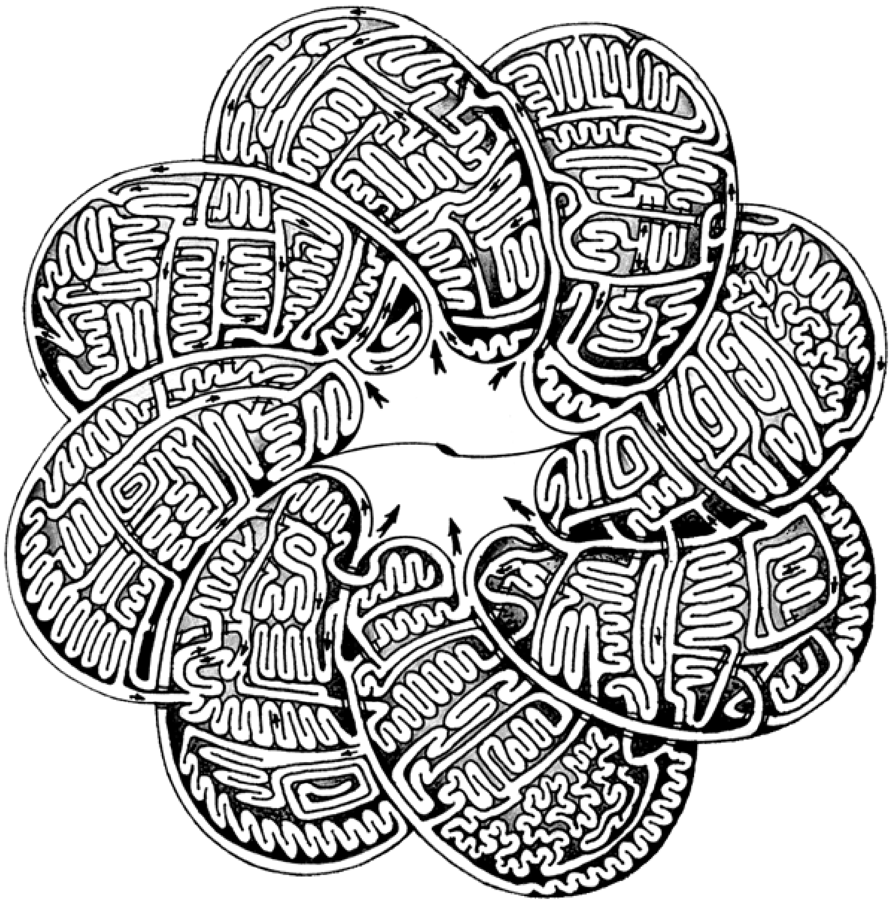

NP

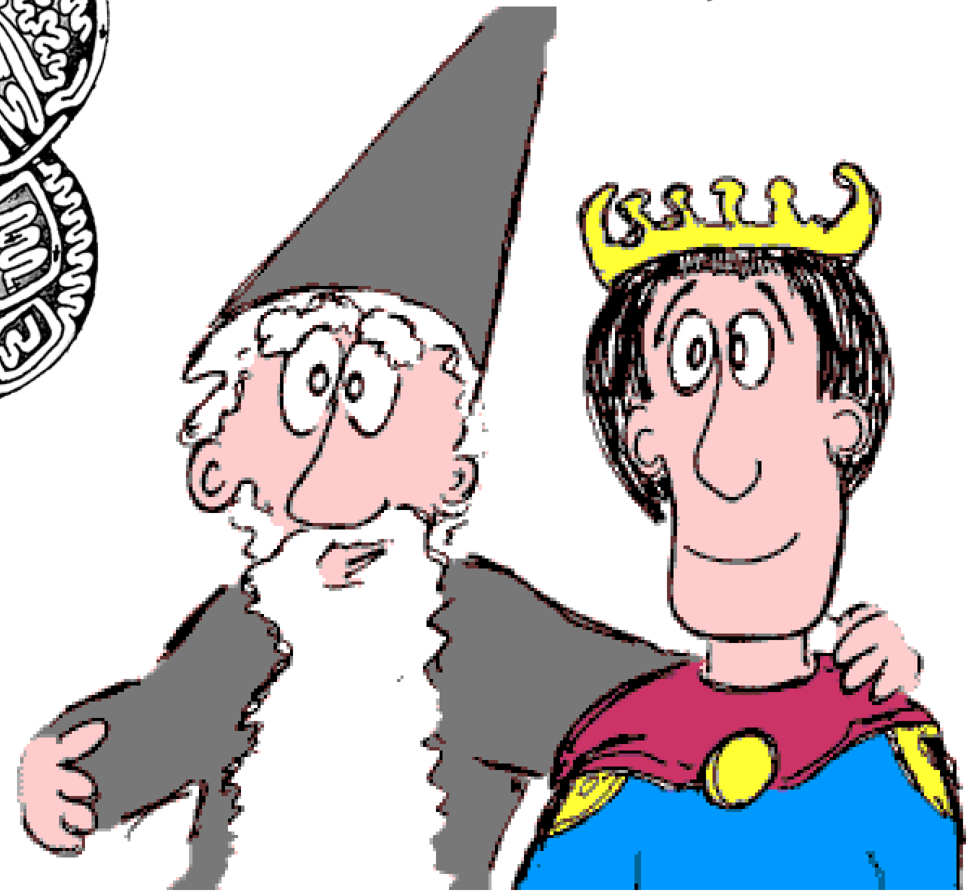

[maze: Andrew Bernhardt]
[A+M: primaryresources.co.uk]

1 A graph isomorphism puzzle

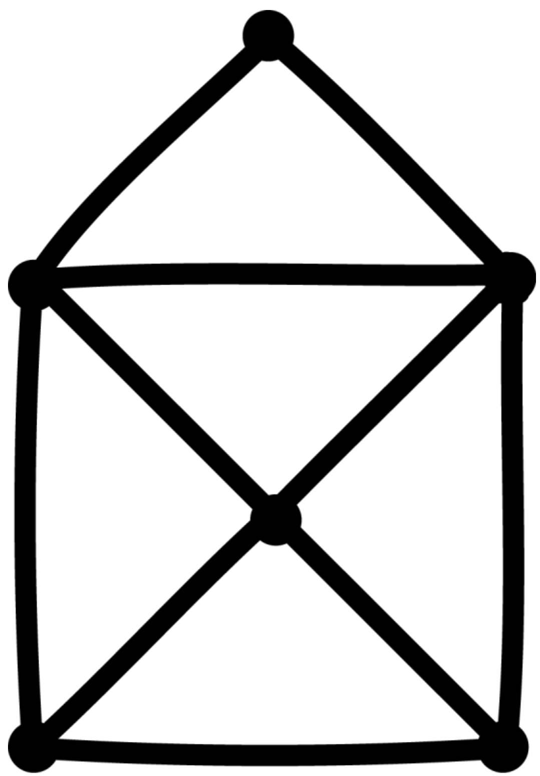

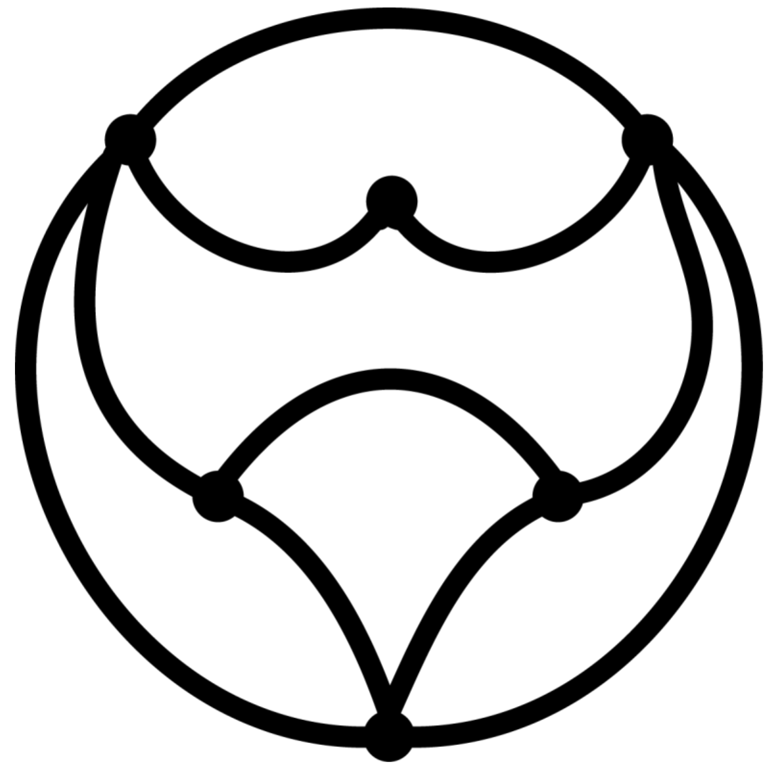

1 A graph isomorphism puzzle

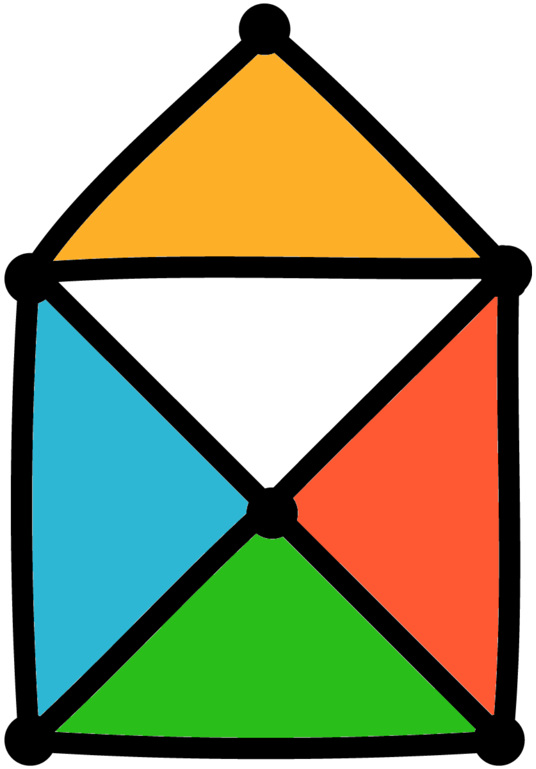

1 A cryptarithmetic puzzle

```
   DID
+DINOS
------
 CROAK
```

?

[www.iread.it/cryptarithms.php]

1 A cryptarithmetic puzzle

$$\begin{array}{r} 595 \\ +\ 59842 \\ \hline 60437 \end{array}$$

!

a proof

1 The NP protocol

Did dinosaurs exist?

a witness

1 The class NP

Yes/no questions, easy to verify solutions.

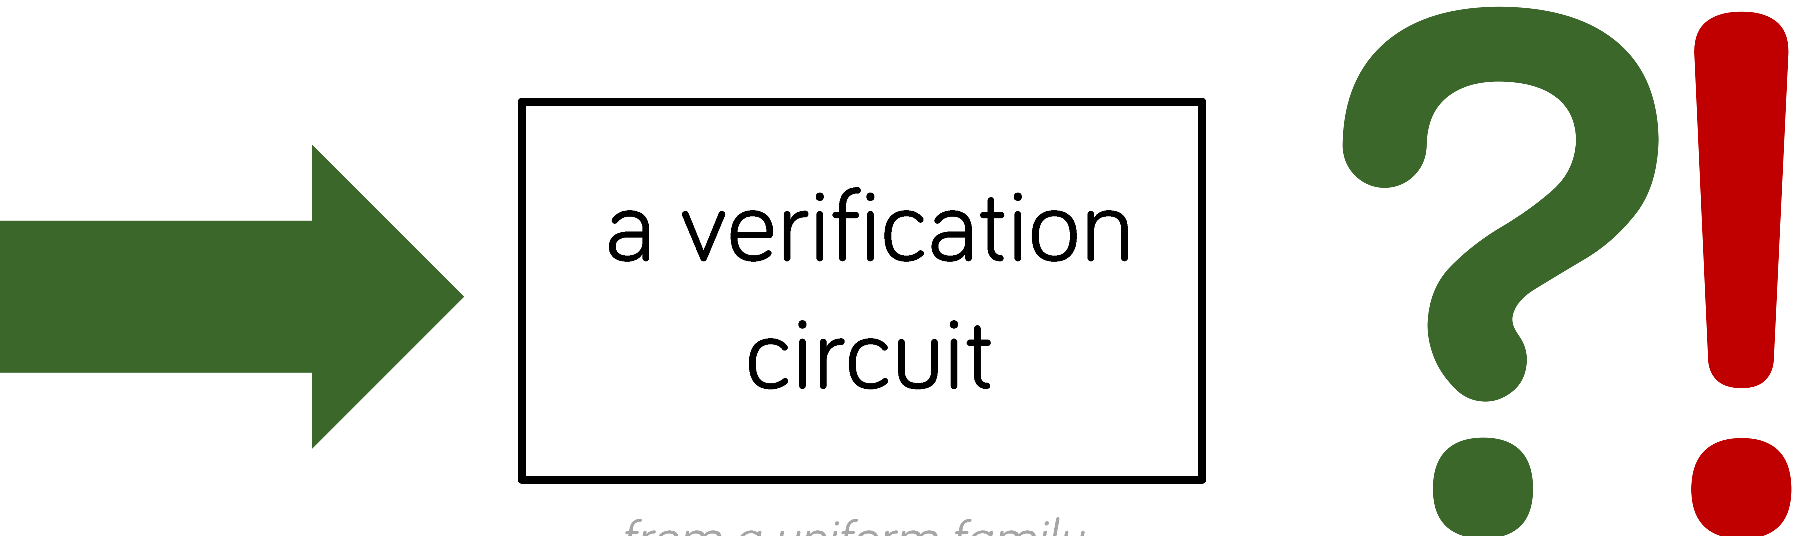

from a uniform family

YES? Accept a good proof.

NO? Reject any witness.

1 The class P

Yes/no questions that we can answer.

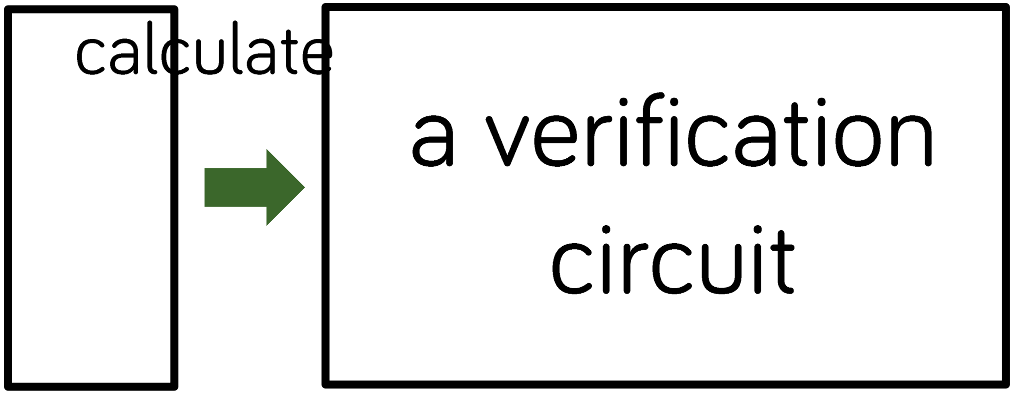

YES? Figure it out by yourself.

NO? Figure it out by yourself.

1 The class NP

Yes/no questions, easy to verify solutions.

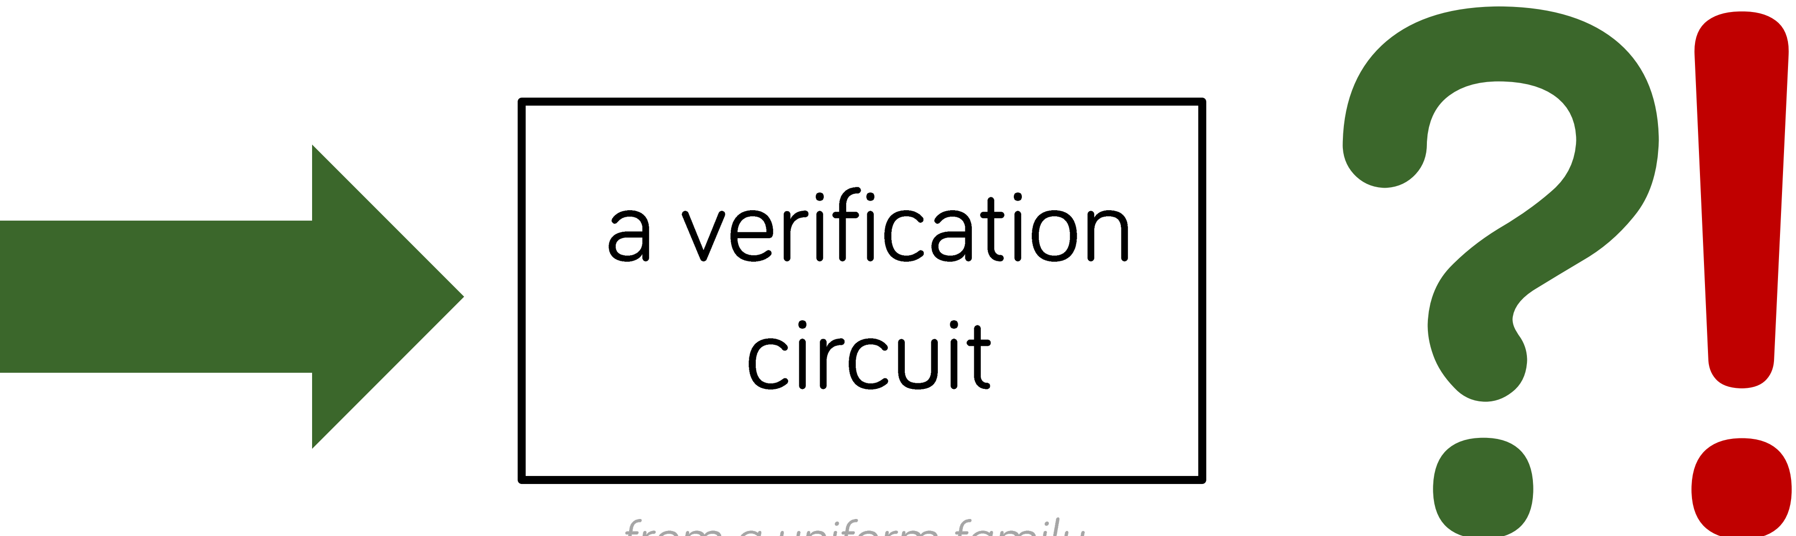

from a uniform family

YES? Accept a good proof.

NO? Reject any witness.

1 An NP-hard problem

- An NP-hard problem solver solves anything in NP.

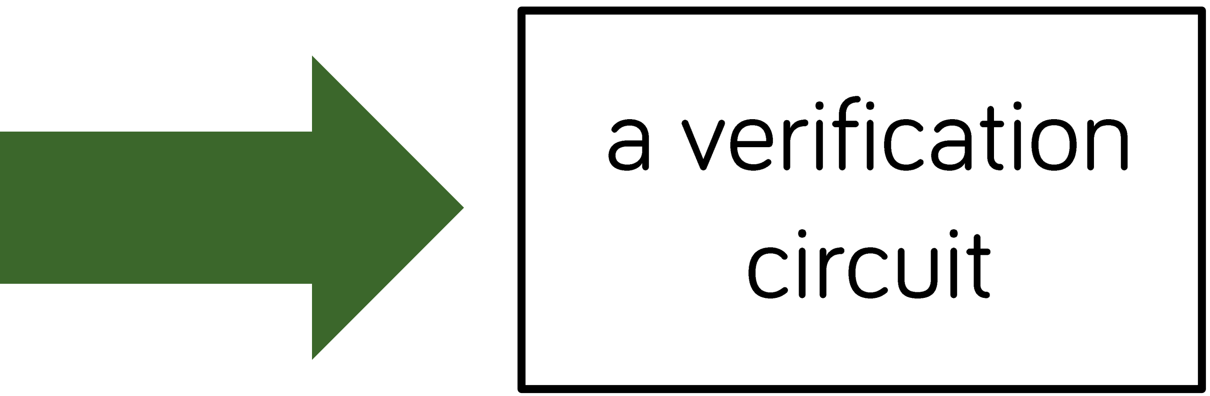

An NP-hard problem

- An NP-hard problem solver solves anything in NP.

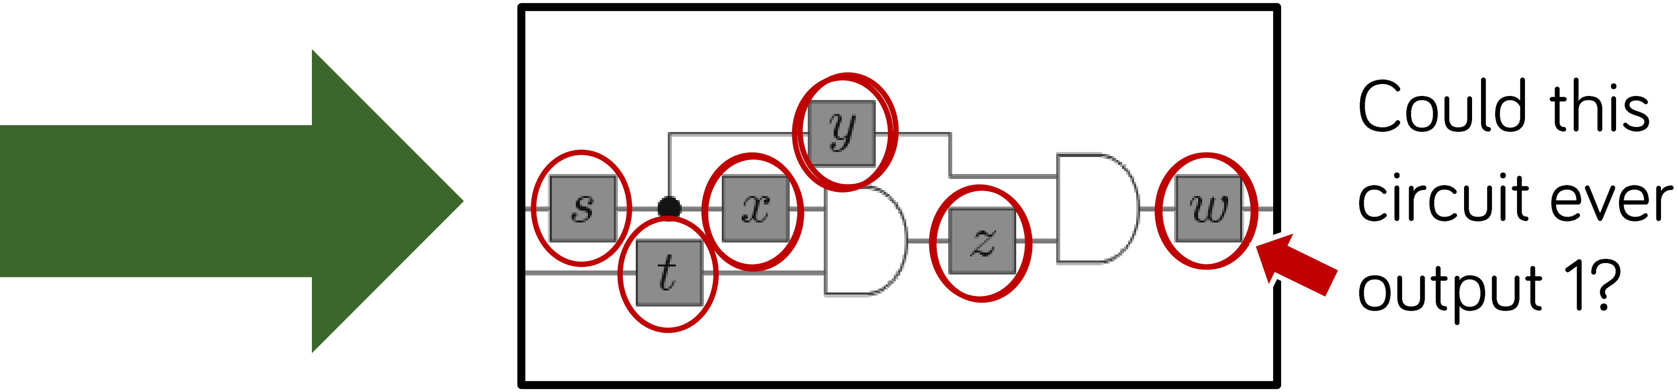

$$(\cdots \vee \cdots \vee \cdots) \wedge (\cdots \vee \cdots \vee \cdots) \wedge \cdots$$

- 3-SAT is NP-hard.
- 3-SAT is in NP.

- An NP-hard problem solver solves anything in NP.

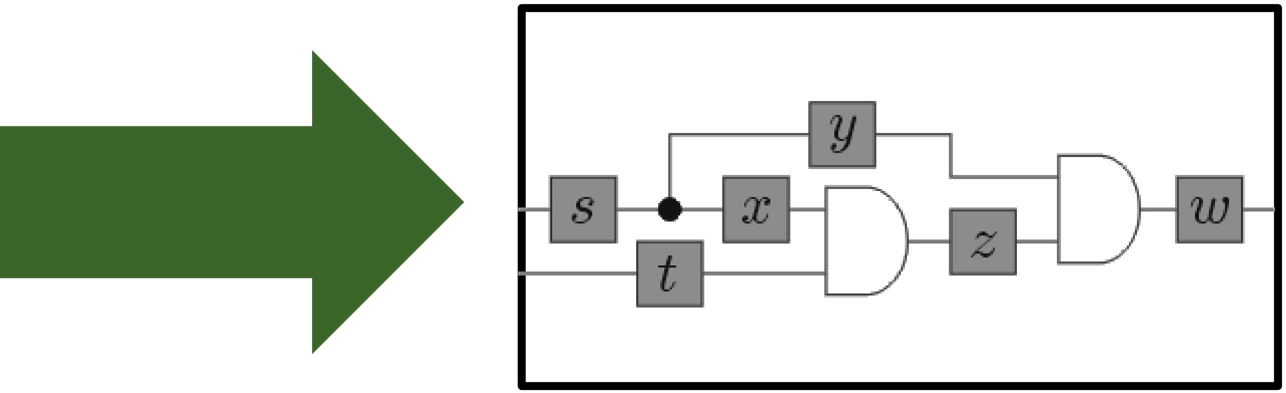

- 3-SAT is NP-complete. [Cook, Levin]

1 NP-complete problems: Hamiltonian cycle

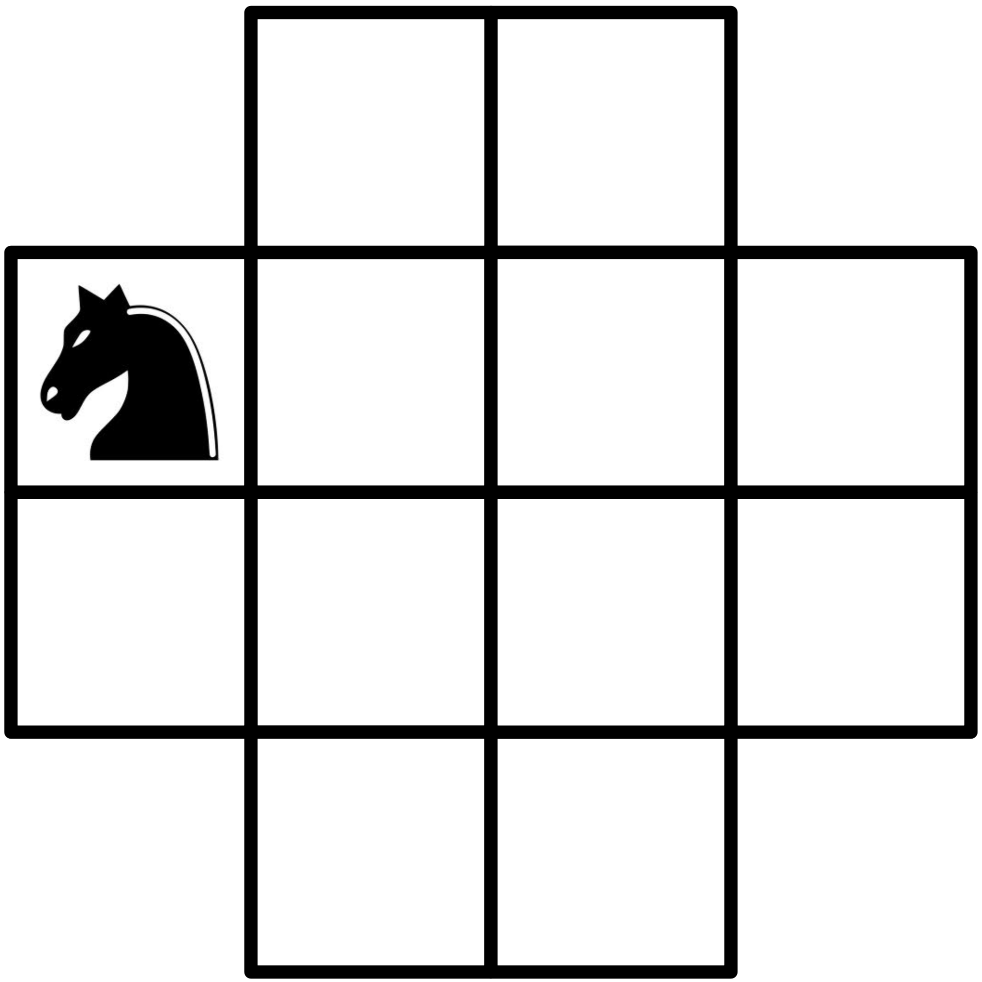

1

NP-complete problems: Hamiltonian cycle

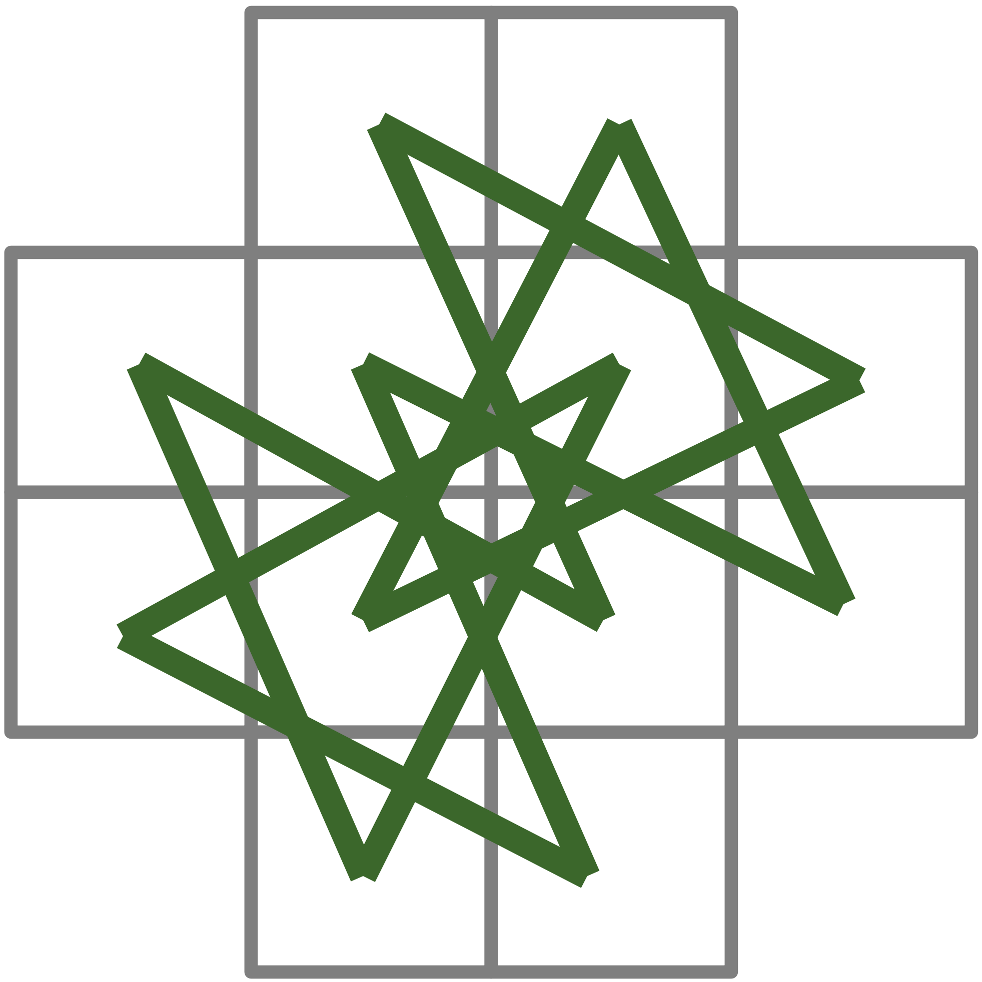

1 NP-complete problems: Ising spin glass

PHYSICAL REVIEW B VOLUME 49, NUMBER 18 1 MAY 1994-II

Ground-state magnetization of Ising spin glasses

Francisco Barahona
Thomas J. Watson Research Center, IBM, Yorktown Heights, New York 10598
(Received 10 December 1993)

We study the magnetization of the ground states of Ising spin glasses as a function of the magnetic field. The grids are two-dimensional with Gaussian bond distribution. Finding a ground state in this case is a difficult combinatorial problem. We present a method that guarantees finding a *true* ground state. Although it is unlikely that one can have an algorithm with a polynomial bound for the running time, in all cases we tried, our method found a solution in a reasonable amount of time. We used this to obtain estimations of the critical exponents.

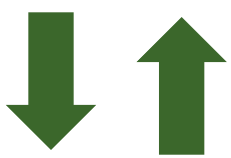

minimize $E = \sum J_{ij} S_i S_j + F \sum S_i$

subject to $S_i \in \{-1, 1\}$.

- using a solver for Max-CUT to find ground states adiabatic

quantum

optimization

2 Adiabatic quantum optimization

[Farhi et al.]

- ground state: minimize a cost function $H_P |z\rangle = h(z) |z\rangle$
- time-dependent, slow-changing $H(t) = \left(1 - \frac{t}{T}\right) H_B + \frac{t}{T} H_P$

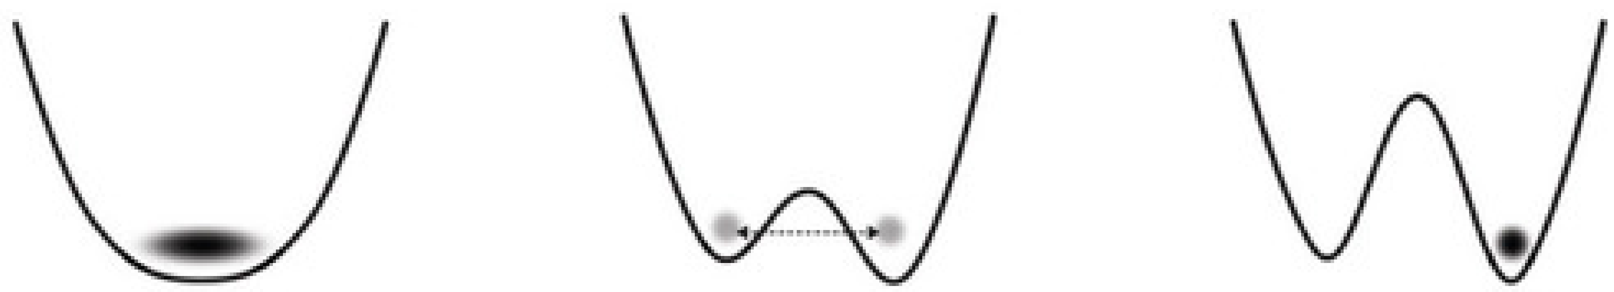

2 Adiabatic quantum optimization

[Farhi et al.]

- ground state: minimize a cost function $H_P \left| z \right\rangle = h(z) \left| z \right\rangle$
- time-dependent, slow-changing $H(t) = \left(1 - \frac{t}{T}\right) H_B + \frac{t}{T} H_P$

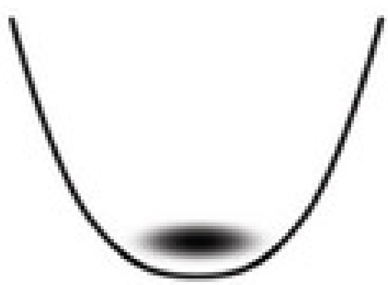

2 Adiabatic quantum optimization

[Farhi et al.]

- ground state: minimize a cost function $H_P \left| z \right\rangle = h(z) \left| z \right\rangle$
- time-dependent, slow-changing $H(t) = \left(1 - \frac{t}{T}\right) H_B + \frac{t}{T} H_P$
- the adiabatic theorem:

Start in an eigenstate?

Stay near an eigenstate.

- how slow is “slow enough”?
- gap scaling with system size
- error correction for AQC???

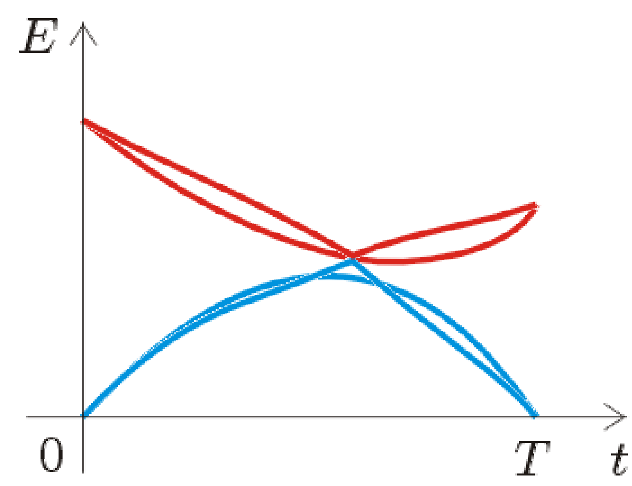

[engadget.com]

D-Wave sells first commercial quantum computer to Lockheed Martin

By Sean Hollister posted May 29th 2011 2:02AM

Who found ten million dollars to drop on the first commercially available quantum computer? Lockheed Martin, it seems, as the aerospace defense contractor has just begun a "multi-year contract" with the quantum annealing experts at D-Wave to develop... nothing that they're ready or willing to publicly discuss at this time.

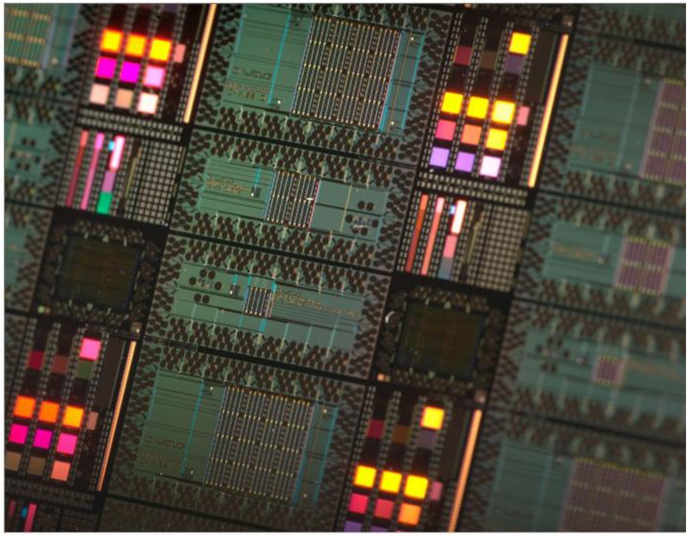

D-Wave's processors on wafer (Courtesy: D-Wave).

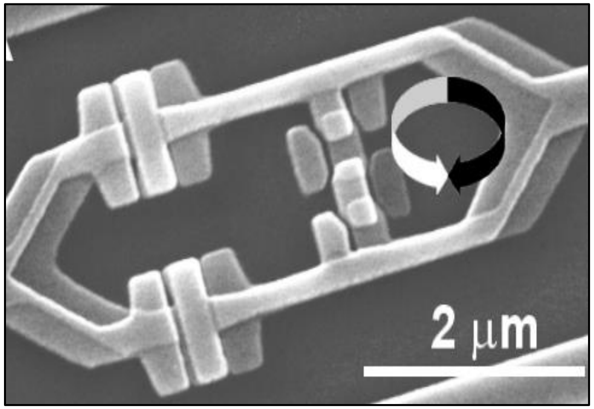

[FSU]

superconducting qubits

D-Wave's Geordie Rose in front of the firm's D-Wave One system (Courtesy: D-Wave).

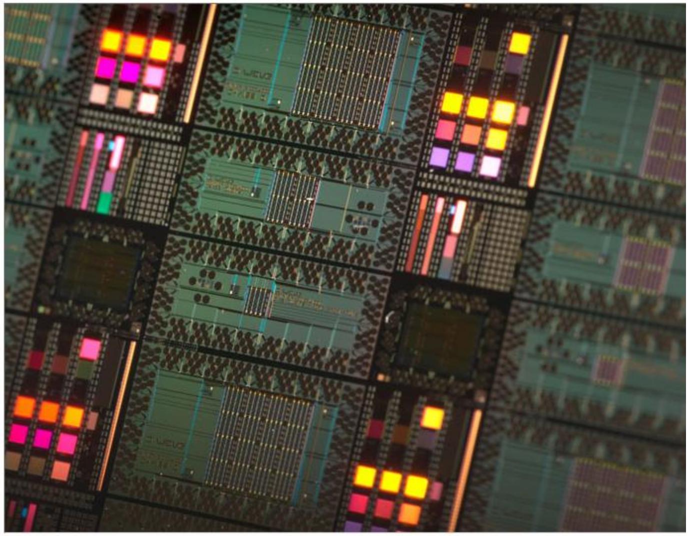

D-Wave's processors on wafer (Courtesy: D-Wave).

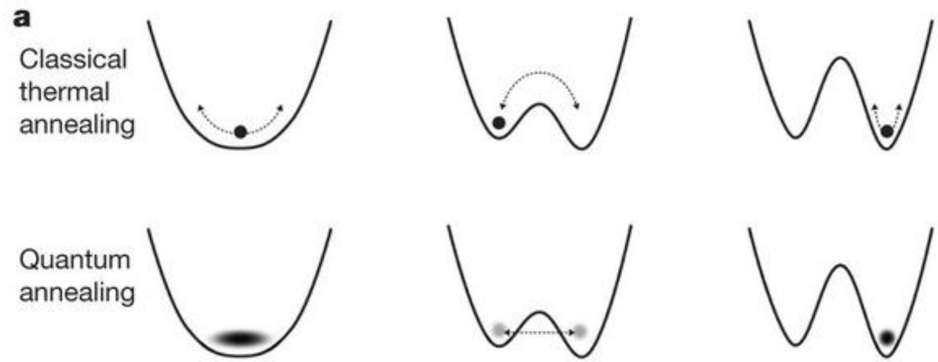

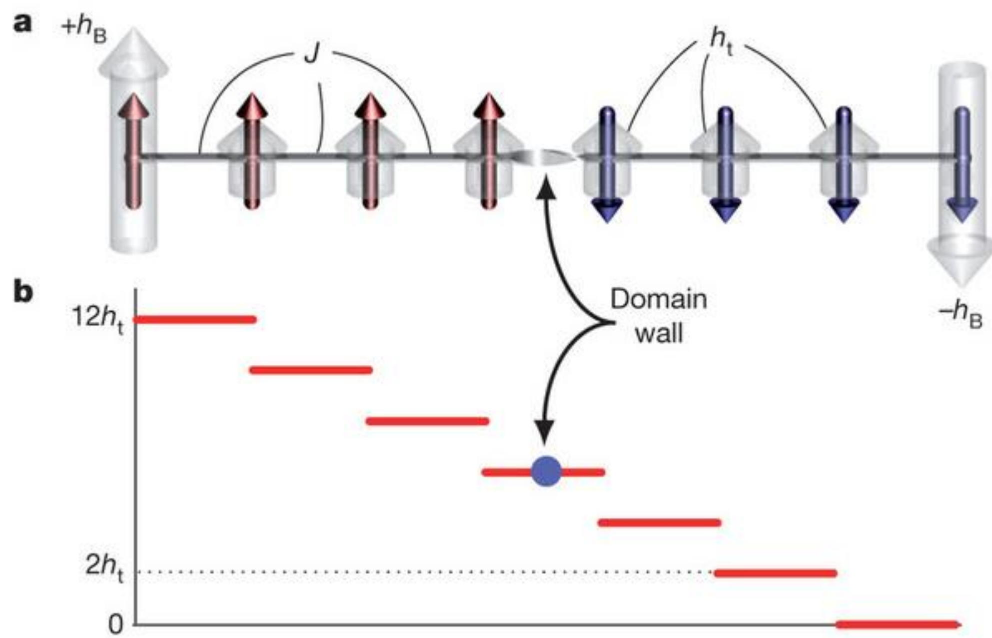

M. W. Johnston et al.

Quantum annealing with manufactured spins

Nature 473, 194–198 (2011)

a frustrated 8-spin chain
confirmed quantum annealing


Shtetl-Optimized
The Blog of Scott Aaronson
Quantum computers are not known to be able
to solve NP-complete problems in polynomial time,
and can be simulated classically with exponential slowdown.
PSPACE
PostBQP
NP
BQP
P


QUANTUM
COMPUTING SINCE
DEMOCRITUS
SCOTT AARONSON

Shtetl-Optimized

The Blog of Scott Aaronson

Quantum computers are not known to be able to solve NP-complete problems in polynomial time, and can be simulated classically with exponential slowdown.

Search Results

« Previous Entries

TIME's cover story on D-Wave: A case study in the conventions of modern journalism

Thursday, February 6th, 2014

Posted in Nerd Interest, Quantum, Speaking Truth to Parallelism | 35 Comments »

D-Wave: Truth finally starts to emerge

Thursday, May 16th, 2013

Posted in Quantum, Speaking Truth to Parallelism | 676 Comments »

My visit to D-Wave: Beyond the roast-beef sandwich

Tuesday, February 21st, 2012

D-Wave Search

You have searched the Shtetl-Optimized weblog

Commercial quantum computer leaves PC in the dust

McGeoch and her colleague Cong Wang of Simon Fraser University, in Burnaby, ran the problem on a D-Wave Two computer, which has 439 qubits formed from superconducting niobium loops. They also tried to solve the problem using three leading algorithms running on a high-end desktop computer. The D-Wave machine turned out to be around 3600 times faster than the best conventional algorithm.

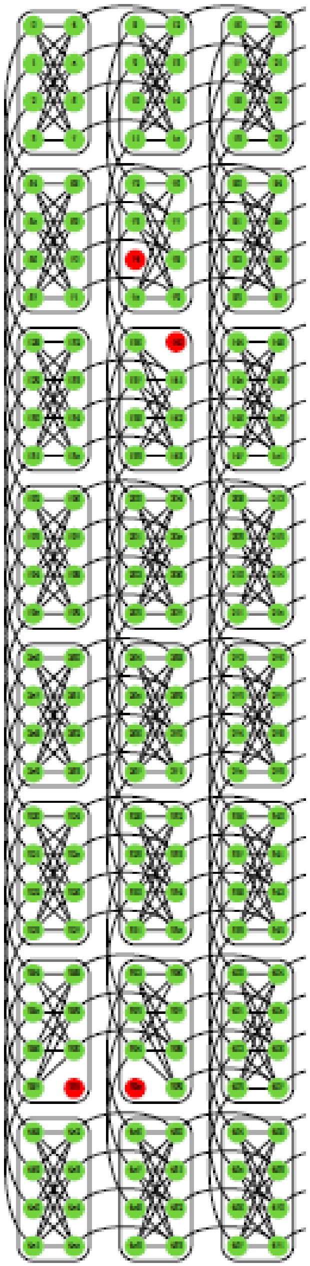

Science 25 July 2014: Vol. 345 no. 6195 pp. 420-424

Defining and detecting quantum speedup

Troels F. Rønnow[1], Zhihui Wang[2,3], Joshua Job[3,4], Sergio Boixo[5,6], Sergei V. Isakov[7], David Wecker[8], John M. Martinis[9], Daniel A. Lidar[2,3,4,6,10], Matthias Troyer[1,*]

+ Author Affiliations

↵*Corresponding author. E-mail: troyer@phys.ethz.ch

ABSTRACT | EDITOR'S SUMMARY

The development of small-scale quantum devices raises the question of how to fairly assess and detect quantum speedup. Here, we show how to define and measure quantum speedup and how to avoid pitfalls that might mask or fake such a speedup. We illustrate our discussion with data from tests run on a D-Wave Two device with up to 503 qubits. By using random spin glass instances as a benchmark, we found no evidence of quantum speedup when the entire data set is considered and obtained inconclusive results when comparing subsets of instances on an instance-by-instance basis. Our results do not rule out the possibility of speedup for other classes of problems and illustrate the subtle nature of the quantum speedup question.

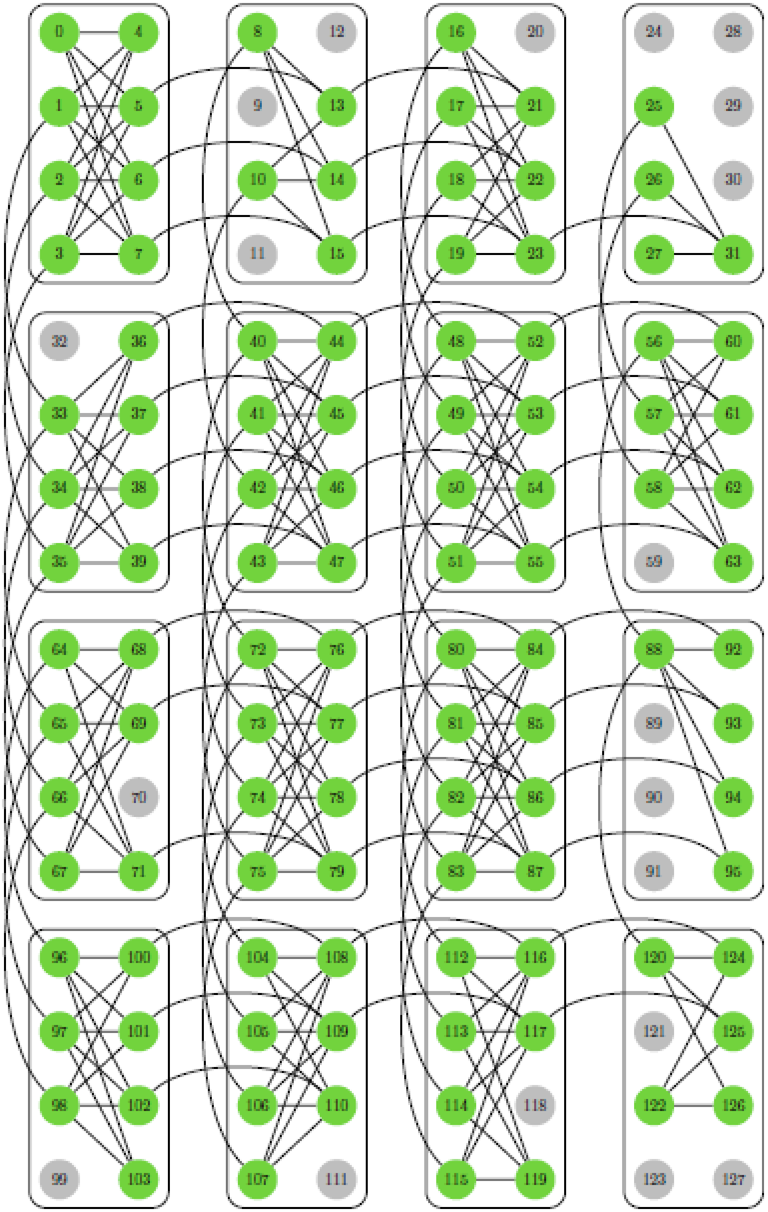

The D-Wave Hamiltonian

external (local) fields

$$-A(t) \sum_{1 \leq i \leq n} \sigma_i^x$$

tunable ferro/AF couplings

$$-B(t) \sum_{i<j} J_{ij} \sigma_i^z \sigma_j^z$$

SCIENTIFIC AMERICAN™

Quantum Chaos: After a Failed Speed Test, the D-Wave Debate Continues

By Seth Fletcher | June 19, 2014 | 9

The views expressed are those of the author and are not necessarily those of Scientific American.

Email Print

Photo courtesy of D-Wave Systems Inc.

How hard can it be to determine whether a computer works as promised? Step one: turn it on. Step two: Try to solve some problems. If it doesn't work, it doesn't work. Right?

Things are never so simple in the real world, of course. And on the highly contested frontiers of quantum computing, matters are more complex still.

DATA CENTRE > HPC

D-Wave disputes benchmark study showing sluggish quantum computer

Faulty benchmarks and sampling methods all wrong, claim Canadian quantumoids

By Iain Thomson, 20 Jun 2014 Follow 2,172 followers

"Eyeballing this treasure trove of data, we're now trying to identify a class of problems for which the current quantum hardware might outperform all known classical solvers," Google said. "But it will take us a bit of time to publish firm conclusions."

Williams said that D-Wave hadn't heard any grumbling from customers after the Troyer paper, and work is continuing as normal. He also said that the firm has a new, upgraded quantum processing system coming out by the end of the year that will increase speeds further. ®

News archive

- 2014
  - September 2014
  - August 2014
  - July 2014
  - June 2014
  - May 2014
  - April 2014
  - March 2014
  - February 2014
  - January 2014
- 2013
- 2012
- 2011
- 2010
- 2009
- 2008
- 2007
- 2006
- 2005
- 2004
- 2003
- 2002
- 2001
- 2000
- 1999

Is D-Wave's quantum computer actually a quantum computer?

Jun 20, 2014 14 comments

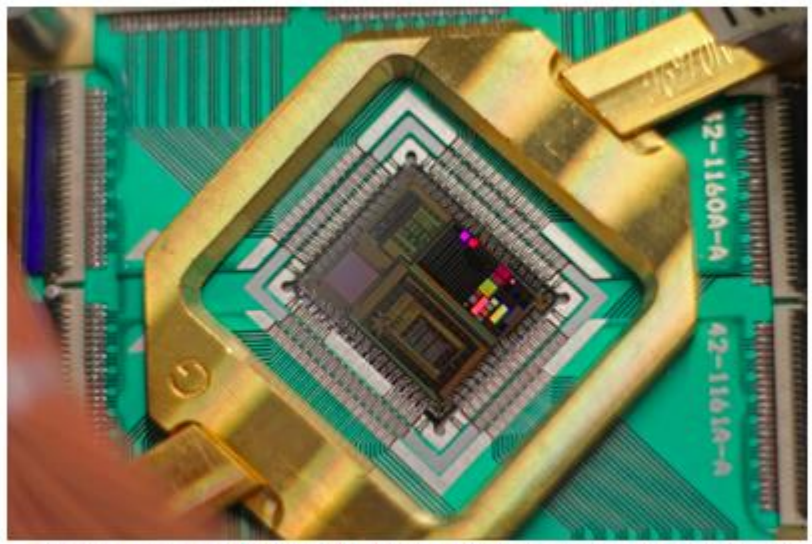

Testing times for D-Wave Systems: the firm's 512-qubit processor

A team of quantum-computing experts in the US and Switzerland has published a paper in *Science* that casts doubt over the ability of the D-Wave Two quantum processor to perform certain computational tasks. The paper, which first appeared as a preprint earlier this year, concludes that the processor – built by the controversial Canadian firm D-Wave Systems – offers no advantage over a conventional computer when it is used to solve a benchmark computing problem.

How "Quantum" is the D-Wave Machine?

Seung Woo Shin[*], Graeme Smith[†], John A. Smolin[†], and Umesh Vazirani[*]

[*] *Computer Science division, UC Berkeley, USA.*
[†] *IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA.*

Abstract

Recently there has been intense interest in claims about the performance of the D-Wave machine. Scientifically the most interesting aspect was the claim in [7], based on extensive experiments, that the D-Wave machine exhibits large-scale quantum behavior. Their conclusion was based on the strong correlation of the input-output behavior of the D-Wave machine with a quantum model called simulated quantum annealing, in contrast to its poor correlation with two classical models: simulated annealing and classical spin dynamics. In this paper, we outline a simple new classical model, and show that on the same data it yields correlations with the D-Wave input-output behavior that are at least as good as those of simulated quantum annealing. Based on these results, we conclude that classical models for the D-Wave machine are not ruled out. Further analysis of the new model provides additional algorithmic insights into the nature of the problems being solved by the D-Wave machine.

How "Quantum" is the D-Wave Machine?

Seung Woo Shin*, Graeme Smith†, John A. Smolin†, and Umesh Vazirani*

3 Results

In this paper, we consider the following classical model: each spin i is modeled by a classical magnet pointing in some direction θ_i in the XZ plane (since there is no σ^y term in our time-dependent Hamiltonian, we can assume that there is no Y component). We assume that the angle θ_i is measured in relation to $\hat{x}$. We further assume that there is an external magnetic field of intensity $A(t)$ pointing in the $\hat{x}$-direction, and that neighboring magnets are coupled via either ferromagnetic or anti-ferromagnetic coupling, according to the specification of the input graph. Our Hamiltonian can then be rewritten as follows:

$$H(t) = -A(t) \sum_{1 \leq i \leq n} \cos\theta_i - B(t) \sum_{i<j} J_{ij} \sin\theta_i \sin\theta_j .$$

Thus, in the absence of noise, i.e. at zero temperature, each spin i will simply align with the net effective field at that location which is given by $A(t)\hat{x} + B(t)\hat{z} \sum_j J_{ij} \sin\theta_j$.

How "Quantum" is the D-Wave Machine?

Seung Woo Shin*, Graeme Smith†, John A. Smolin†, and Umesh Vazirani*

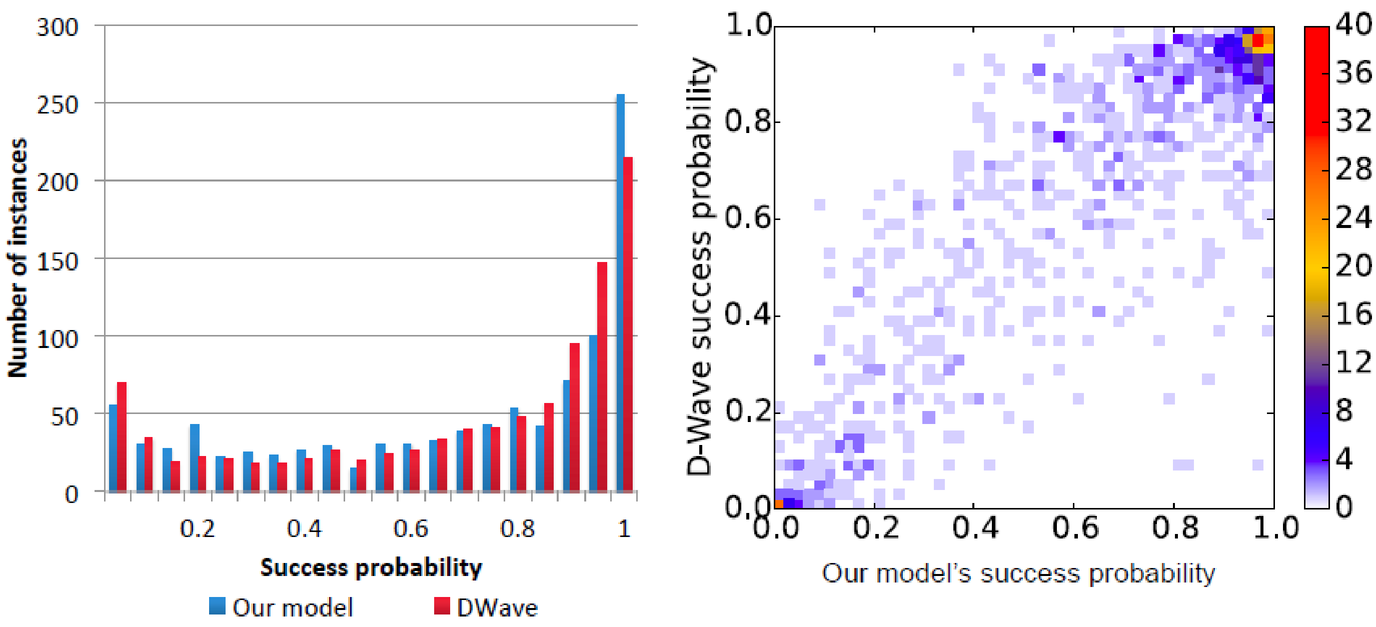

Figure 5: Histogram and scatterplot of our classical model. The correlation coefficient R between the D-Wave One and our model is about 0.89.

MIT Technology Review, 9/3/2014

Google Launches Effort to Build Its Own Quantum Computer

Google's crack at a quantum computer is a bid to change computing forever.

By Tom Simonite on September 3, 2014

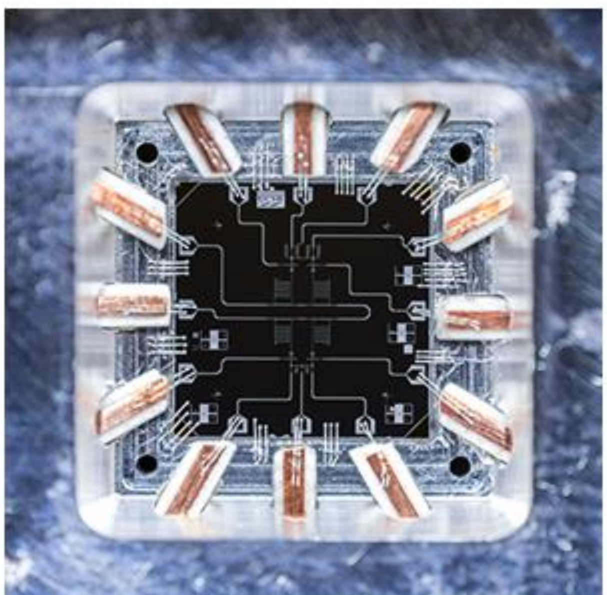

Google is about to begin designing and building hardware for a quantum computer, a type of machine that can exploit quantum physics to solve problems that would take a conventional computer millions of years.

Since 2009, Google has been working with controversial startup D-Wave Systems, which claims to make "the first commercial quantum computer." And last year Google purchased one of D-Wave's machines. But independent tests published earlier this year found no evidence that D-Wave's computer uses quantum physics to solve problems more efficiently than a conventional machine.

WHY IT MATTERS

Raise a martini glass for Google and Martinis!

September 6th, 2014

We've already been discussing this in the comments section of my previous post, but a few people emailed me to ask when I'd devote a separate blog post to the news.

OK, so for those who haven't yet heard: this week Google's Quantum AI Lab announced that it's teaming up with John Martinis, of the University of California, Santa Barbara, to accelerate the Martinis group's already-amazing efforts in superconducting quantum computing. (See here for the *MIT Tech*'s article, here for *Wired*'s, and here for the *WSJ*'s.) Besides building some of the best (if not *the* best) superconducting qubits in the world, Martinis, along with Matthias Troyer, was also one of the coauthors of two important papers that found no evidence for any speedup in the D-Wave machines. (However, in addition to working with the Martinis group, Google says it will *also* continue its partnership with D-Wave, in an apparent effort to keep reality more soap-operatically interesting than any hypothetical scenario one could make up on a blog.)

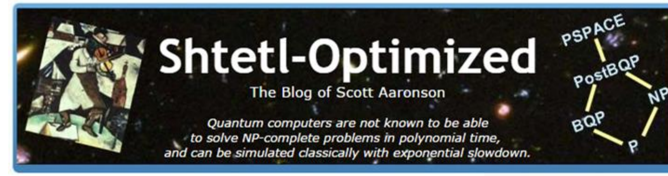

John Martinis

http://web.physics.ucsb.edu/~martinisgroup/