Novel ways of using a quantum annealer

Computability in Europe 2019

Nicholas Chancellor

18 July, 2019

Joint Quantum Centre

Durham University

EPSRC
Pioneering research and skills

Two different approaches to quantum computing

'Gate' based quantum computing

- Discrete quantum operations on qubits
- Construct 'circuits' out of these gates
- Detect and correct errors to reduce effect of noise

Quantum annealing

- Map optimization problem directly to energies of different states
- Allow quantum physics to help search solution space
- Low temperature environment helps solve problems

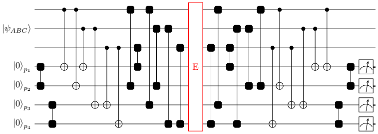

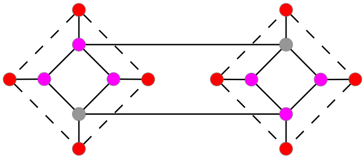

(some) Advantages and disadvantages of each

'Gate' based quantum computing

- Can simulate arbitrary quantum systems
- Error correction can get rid of all noise *in principle*
- Could simulate quantum annealing *in principle*

- Harder to build, largest device is tens of qubits
- All noise likely to be harmful rather than beneficial

Quantum annealing

- Easier to build, largest device is thousands of qubits
- Tolerant to noise, in fact noise helps solve problems
- Naturally produces thermal distributions

- Unclear if error correction is feasible
- Cannot be used for some quantum algorithms as implemented

D-Wave Quantum annealing hardware

- Superconducting circuit devices with up to 2,048 qubits in 16x16 'chimera' configuration
- Operates in a cryostat at ≈ 0.015 K (200x colder than interstellar space: 3K)

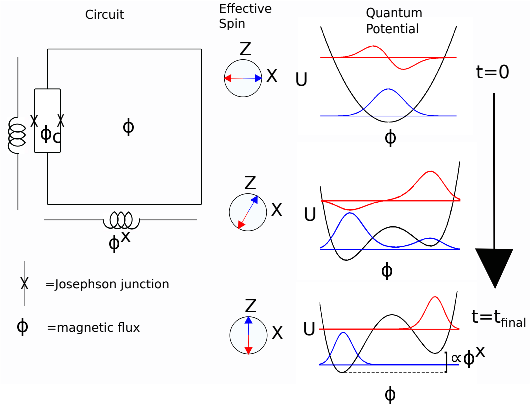

Thermal and quantum fluctuations work together to solve problems

Obligatory slide: D-Wave controversy

Two separate controversies:

1) Are the dynamics actually quantum? Yes!

- Lots of evidence, most striking is simulation of extremely quantum KT phase transition Nature 560 456–460 (2018)
- Classical models reproduce **some** behaviours, expected → mean field approximation

2) Can it ~~beat~~ improve classical computing? Open question

- No conclusive speedup demonstrated yet
- Not what this talk is about

- Currently the only large scale device to study algorithmic application of quantum mechanics
- **Good science can be done regardless of answer to question 2!**

How to actually solve problems with these devices: Optimization (traditional approach)

1. Map problem to one and two body terms of the appropriate form (Ising model) Optimality of solution → energy
2. Embed in hardware graph by strongly linking qubits together to form 'logical' qubits (3x3 chimera shown below)

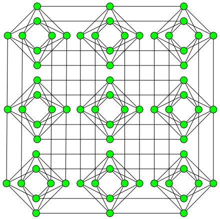

3. Quantum dynamics finds low energy states, run many times and take lowest energy solution

Each run is independent and starts from equal superposition 'state of maximal ignorance' could do better by using information from previous runs (more on this later)

Problem mapping example: maximum independent set

Have:

- Binary variables $Z_i \in \{-1, 1\}$
- Minimisation over Hamiltonian made of single and pairwise terms $\mathrm{H}_{\mathrm{Ising}} = \sum_i h_i Z_i + \sum_{j>i} J_{i,j} Z_i Z_j$

Want:

- Maximum* independent set: how many vertexes on a graph can we colour so none touch? $\rightarrow$ NP hard

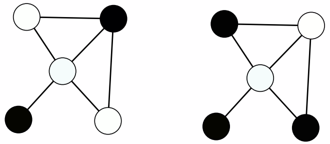

Method:

1. For an edge between vertex i and j add $Z_i + Z_j + Z_i Z_j \rightarrow$ penalizes colouring ($Z = 1$) adacent vertexes
2. Add $-\lambda Z_i$ to reward coloured vertexes ($0 < \lambda < 1$)

*Not to be confused with *maximal* independent set, which is not a hard problem

Minor embedding

- Strong 'ferromagnetic' ($-Z_i Z_j$) coupling energetically penalizes variables disagreeing
- If strong enough than entire 'chain' acts as a single variable
- Mathematically corresponds to mapping one graph to graph minors of another

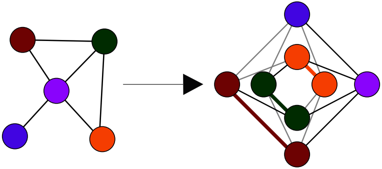

Can embed arbitrary graphs into the hardware graph with polynomial (n^2 for fully connected) overhead → Ising model restricted to hardware graph is also NP-hard

Novel way to use a quantum annealer I: thermal sampling

Systems at finite temperatures ($T = \frac{1}{\beta}$) naturally tend toward a 'Boltzmann' distribution, probability of bitstring $b_i \in \{0, 1\}^n$ with energy $E_i = \langle H_{\text{Ising}} \rangle_i$ is:

$$p(b_i) = \frac{\exp(-\beta E_i)}{\sum_j \exp(-\beta E_j)}$$

- βE acts as log-probability
- Each element of H_{Ising} corresponds to probability of a condition being met
- Can conditionally sample a distribution

Example: decoding of (classical) communications

- Data and parity checks both transmitted on noisy channel
- Want to conditionally sample: most likely series of errors *given* parity check results

Maximum entropy versus maximum likelihood

Two ways to decode communications:

1) Maximum likelihood (ML)

- Correct single most likely series of errors
- Optimization problem: maximize probability

2) Maximum entropy (ME)

- Conditionally sample probability distribution $\rightarrow$ requires error rate to be known at least approximately
- Take a 'vote' if a correction helps in more cases than it hurts, do it

ME is always as good or better than ML, but harder to do

Can a quantum annealer doing approximate ME beat perfect ML?

Proof-of-concept experiments*

- Use small code which matches D-Wave hardware graph
- Perform exact ML decoding using Bucket tree elimination
- Perform approximate ME using D-Wave device $\rightarrow$ scale problem by α to control effective temperature
- Compare, see which wins at different error rates

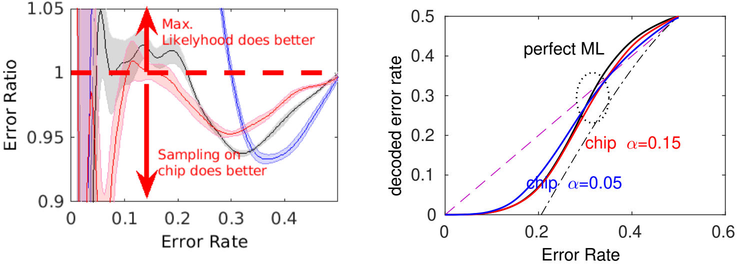

Over some range, ME on real chip beats perfect ML!

see: Scientific Reports vol. 6, 22318 (2016) for details

*work done at UCL, supported by EPSRC and Lockheed Martin

The bigger picture for sampling

Beyond proof-of-concept (ongoing project):

- Embedded and encoded problems, can sampling still work
- When is the device able to approximately sample thermally?
- Hybrid quantum/classical algorithms to improve sampling

Boltzmann machines?

- Neural networks where sampling thermal distribution is required to train
- Area of active research

Conditional sampling problems beyond (classical) decoding

- Can be extended to quantum error correction: arχiv:1903.10254
- Many other areas where conditional sampling may be useful...

Novel way to use a quantum annealer II: reverse annealing

Hybrid (quantum/classical) algorithms

- Many good classical optimization algorithms already exist
- Need every advantage we can get to maximally use early quantum hardware

Why reverse annealing?

- Easy inclusion of previously found solutions in algorithm calls (search range controlled by parameter s': $s' = 1 \rightarrow$ no search,$s' = 0 \rightarrow$ traditional annealing)
- Flexible: can be used with most existing techniques
- Now available on D-Wave devices

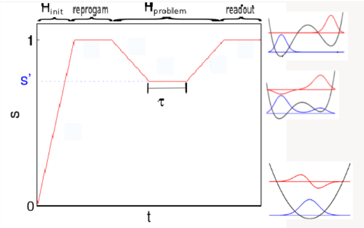

Cartoon example: energy landscape with rough and smooth features (see: NJP 19, 2, 023024 (2017))

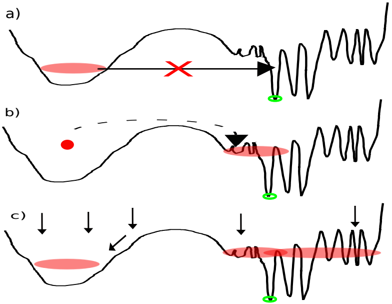

a) QA gets stuck in broad local minima and cannot tunnel to correct minima
b) Classical algorithms can easily explore the broad features, while the annealer can explore the rough ones
c) Even random initialization can improve solution probabilities, may hit rough region by chance

Proof-of-principle experiments

Construct a problem Hamiltonian with the following properties:

1. Wide false energy minimum which 'tricks' traditional quantum annealing algorithm
2. Relatively narrow true minimum energy
3. Local minimum near true minimum for start state

Joint work with Viv Kendon, funded by NQIT and EPSRC

NQIT Networked Quantum Information Technologies

EPSRC Pioneering research and skills

Experimental results

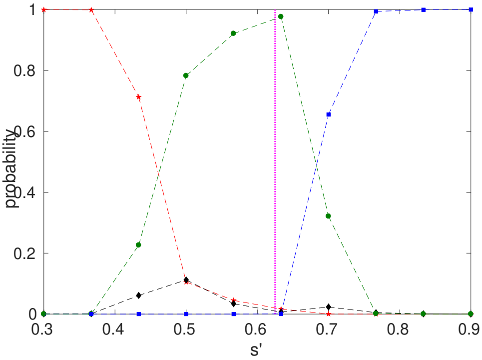

- Level crossing between true ground state and false minima at magenta line → energetically preferable to be in narrow minimum to right of line and broad to the left
- Anneal at maximum allowed rate, wait time (τ) of $20\mu S$
- Frozen in starting state for small s', find true minimum at moderate s', trapped for large s'

Reverse annealing in algorithms*

1. Start from one ground state to find other ground states (D-Wave whitpaper 14-1018A-A†)
   - Finding other GS 150x more likely then forward
2. Search locally around classical solution (arχiv:1810.08584†)
   - Start from greedy search solution
   - Speedup of 100x over forward annealing
3. Iterative search (arχiv:1808.08721†)
   - Iteratively increase search range until new solution found
   - Forward annealing could not solve any, reverse solved most
4. Quantum simulation (Nature 560 456–460 (2018)†)
   - Seed next call with result from previous
   - Seeding with previous state makes simulation possible
5. Genetic algorithms (arχiv:1907.00707†)
   - Used reverse annealing for mutations
   - Out-performed state-of-the-art solvers by orders of magnitude on some test cases
6. Proposals for Monte Carlo and Genetic like algorithms (NJP 19, 2, 023024 (2017) and arχiv:1609.05875)

*† indicates experimental results

Novel way III: solution robustness*

Using quantum annealers to find good solutions near other good solutions

- Already known that annealers preferentially find good solutions which are 'near' other good solutions → leverage these effects algorithmically
- If a good solution is already known, can we use an annealer to trade optimality for robustness?

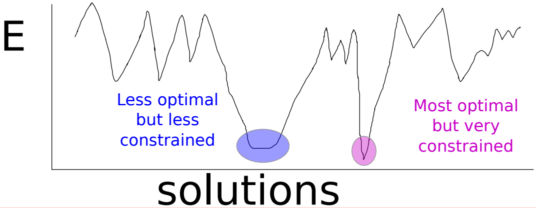

Funded by BP, NQIT, and EPSRC, work with Simon Benjamin group

bp

NQIT Networked Quantum Information Technologies

EPSRC Pioneering research and skills

*outline of results, not time for full presentation, see my AQC 2019 or BCTCS talk (nicholas-chancellor.me/presentations) or ask me for full version

A simple (motivational) example

Consider 16 qubit gadget from N. G. Dickson et. al. Nature Comm. 4, 1903 (2013) :

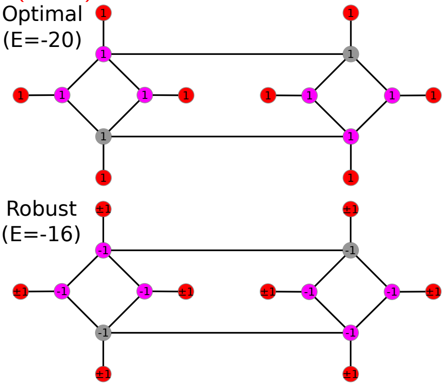

- a is the ground state but
- A D-Wave 2000Q with $1,280,000$ $5\mu s$ runs finds b $1,277,824$ times and a only 17 times

Outline of experiments and results

1. Set up problem with known solution
2. Add 'gadgets' which can increase frustration for an energy cost
3. Reverse annealing starting from optimal solution, find robust solutions

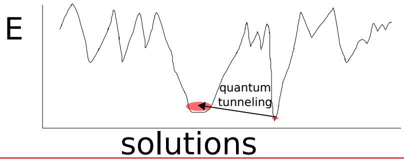

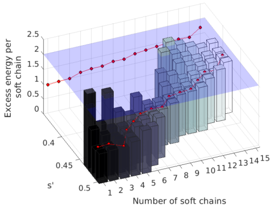

results:

- Non-trivial tradeoff can be performed
- Can find better solutions than without reverse annealing

Take home messages

D-Wave quantum annealers

- Opportunity to do experimental CS on large quantum systems $\rightarrow$ much more experimentally mature than gate model
- No conclusive speedup yet, but science can be useful regardless

Thermal sampling:

- Physical systems naturally tend to thermal distribution, useful for conditional sampling
- Proof-of-principle experiment shows this can work

Reverse annealing:

- Ability to 'seed' known good solution allows many algorithmic possibilities
- Can be used to trade of optimality for robustness of solutions