Research Spotlight Forum

6.4.19 Quantum Information Sciences

Lessons from Quantum Annealing for the NISQ Era

PRESENTED BY:

Tameem Albash

Electrical and Computer Engineering, CQuIC, @UNM (Fall 2019)

talbash@unm.edu

Sandia National Laboratories | Georgia Institute of Technology | Illinois | Purdue University | The University of Texas at Austin | The University of New Mexico | NM State University | New Mexico Tech Science • Engineering • Research University

SAND2018-14303 O

SUMMARY OF MYSELF

- Currently at ISI@USC but will be moving to UNM in the fall.
- I started in QIS as part of the USC team when Lockheed-Martin installed the first D-Wave quantum annealing processor at ISI, where I played a role in developing open-quantum-system simulations and benchmarking tools to detect signatures of 'quantumness' on these devices. Currently part of the performer team of IARPA's Quantum Enhanced Optimization (QEO) program.
- While much of my focus to date has been on the adiabatic paradigm of quantum computing, I hope to leverage many of the important lessons we have learned in the adiabatic approach to explore various aspects of the performance of quantum algorithms in the NISQ era.

Keywords:

Adiabatic quantum computing, variational quantum algorithms, quantum simulation

FUNDING SOURCES

- IARPA: Quantum Enhanced Optimization (QEO)
  - "QEO seeks to harness quantum effects required to enhance quantum annealing solutions to hard combinatorial optimization problems. The physics underlying quantum enhancement will be corroborated by design and demonstration of research-scale annealing test beds comprised of novel superconducting qubits, architectures, and operating procedures. All work will serve to demonstrate a plausible path to enhancement and a basis for design of application-scale quantum annealers."
- Performer team is a multi-university, government, and industry collaboration:
  - USC, USC-ISI, Berkeley, MIT, Harvard, Waterloo, UCL, Tokyo Tech, Saarland, UNM, Northrop Grumman, Lockheed Martin, MIT Lincoln Labs, NASA Ames, Texas A&M

CURRENT WORK IN QUANTUM INFORMATION SCIENCES: TALES FROM THE TRENCHES OF QUANTUM ANNEALING

- Why did we think quantum annealing was the easiest path for demonstrating the usefulness of specialized quantum computing
- Summary of some of the many problems we encountered along the way
- What excites me and opportunities moving forward

WHY MIGHT QUANTUM ANNEALING BE A PROMISING APPROACH

- An algorithm from the adiabatic paradigm quantum computing (1), whose sole purpose is to solve for the ground state of a 'classical' (diagonal in the computational basis) Hamiltonian

$$H(s) = -(1-s)\sum_i \sigma_i^x + s\,H_p, \qquad s = t/t_f \in [0,1]$$

- Adiabatic evolution guarantees high success. Only requires slow anneal of all qubits.
- Only requires measurements when the Hamiltonian is classical (at $s = 1$).

} Minimal control requirements

- Ground state evolution
  - thermal relaxation can only help.
  - dephasing in instantaneous energy eigenbasis is fine.
- Perturbations to the annealing path are often okay (often referred to as unitary control errors). We will however see problems with this later.
- H_p is often taken to be an Ising Hamiltonian. Many combinatorial optimization problems can be easily cast as Ising Hamiltonians (2), so there are a lot of applications that can be easily mapped to it.

(1) E. Farhi, et al., arXiv:quant-ph/0001106 (2000).
(2) F. Barahona, J. Phys. A: Math. Gen. 15 3241 (1982)

BE VERY SKEPTICAL OF SMALL PROBLEMS

- Our ability to classically simulate quantum evolution restricts us to small system sizes. Typical small problems (~ 20 qubits) tend to be artificially easy.
  - Studying small system sizes may lead you to believe that the scaling of the algorithm is polynomial (1), when in fact it transitions to an exponential at larger sizes (2).
- At small sizes, running adiabatically may not be optimal. Fast repeated anneals more advantageous.

Algorithmic performance on small problems may not be representative of performance at larger problem sizes!

Most problems are not amenable to any analytical analysis.
How do we determine the algorithm's performance?
Seems like we are stuck with building the device in order to find out.

(1) E. Farhi, et al., Science 292 (5516), 472 (2001).
(2) Young, A. P., S. Knysh, and V. N. Smelyanskiy, Phys. Rev. Lett. 101, 170503 (2008), Phys. Rev. Lett. 104, 020502 (2010).

THE PROBLEM OF TEMPERATURE

- Even under most innocuous decoherence model (weak coupling limit master equation), temperature remains a severe impediment to performance.
  - At too high temperatures, thermal hopping dominates over quantum tunneling as the mechanism of overcoming energy barriers, resulting in measurable differences in performance. To 'observe' this (1):
    - Identify instances for which a semiclassical limit (spin-vector Monte Carlo) encounters discontinuous jumps in the global minimum of the energy landscape. An adiabatic evolution would tunnel through this barrier.
    - Perform quantum Monte Carlo simulations at high and low temperatures, in order to favor thermal hopping over the energy barrier versus tunneling through the energy barrier.

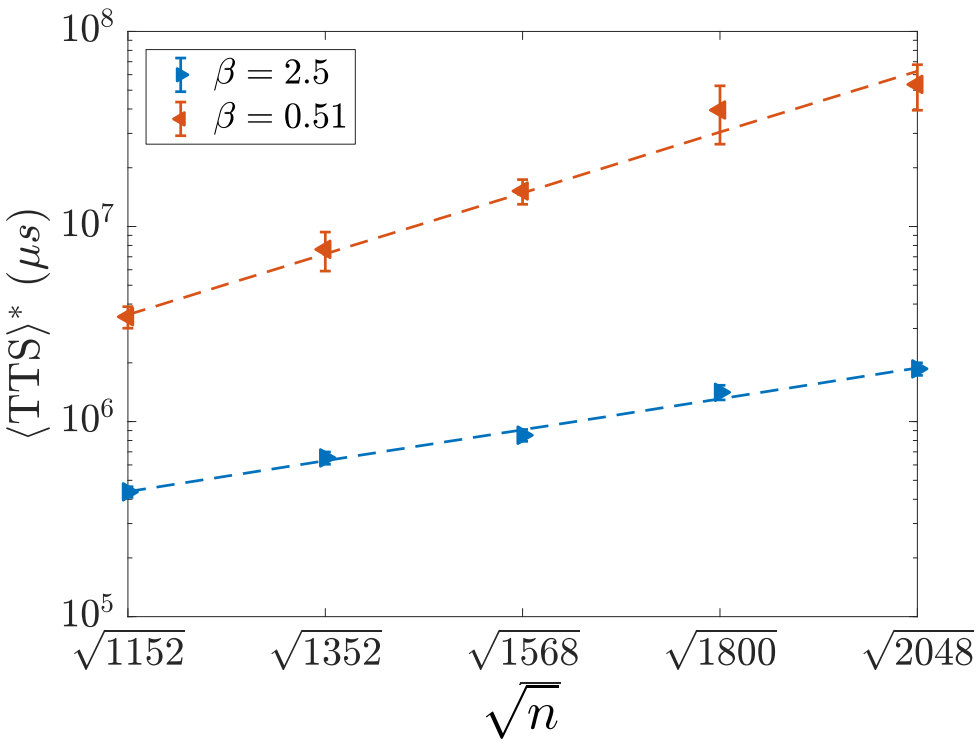

(1) TA and D. A. Lidar, Phys. Rev. X 8, 031016 (2018).

THE PROBLEM OF TEMPERATURE

- Even under most innocuous decoherence model (weak coupling limit master equation), temperature remains a severe impediment to scalability.
  - In with a thermal environment for a long time means system approaches the Gibbs state (1) and not the ground state. Probability of finding the ground state now depends on weight of the ground state in the Gibbs state.

$$H_R = H_P - E_0$$

Residual energy measures distance (in energy) from the ground state. Residual energy distribution for the Gibbs state becomes more Gaussian like for increasing system size at a fixed temperature.

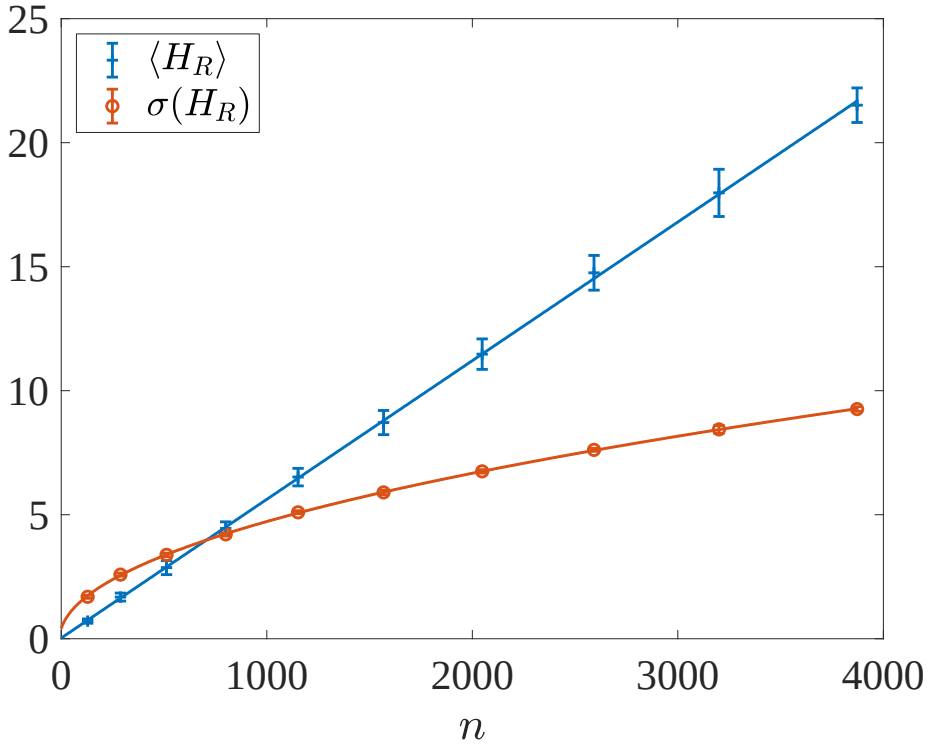

For Gibbs states with residual energy distributions whose mean grows as n and whose standard deviation grows as $\sqrt{n}$, it becomes exponentially unlikely to sample the ground state (2)

(1) L. C. Venuti, TA, D. A. Lidar, and P. Zanardi, Phys. Rev. A 93, 032118 (2016)
(2) TA, V. Martin-Mayor, I. Hen, Phys. Rev. Lett. 119, 110502 (2017).

THE PROBLEM OF IMPLEMENTATION ERRORS

- Implementation of final Hamiltonian is critical. Small deviations in the final Hamiltonian results in a different ground state.
  - Adiabatic evolution, even if ideal, results in the wrong solution!

$$H_P \rightarrow \sum_i (h_i + \delta h_i)\sigma_i^z + \sum_{\langle i,j \rangle} (J_{ij} + \delta J_{ij})\sigma_i^z \sigma_j^z$$

Median probability that the ground state remains unchanged when introducing Gaussian noise $(\mu = 0, \sigma^2)$ on the Ising parameters as the problem size is scaled.

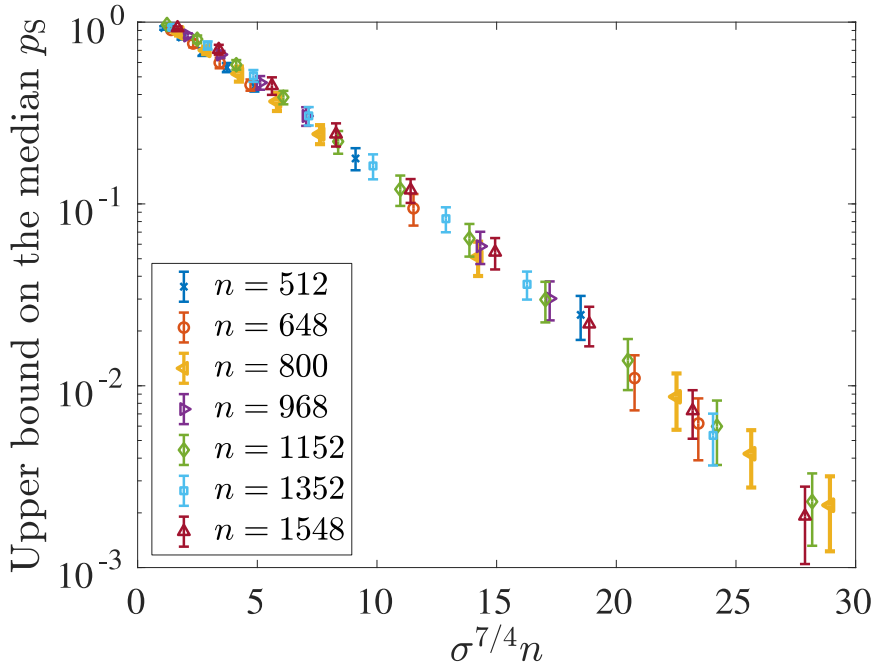

Probability that ground state remains unchanged scales down exponentially in the system size for a fixed noise level σ

(1) TA, V. Martin-Mayor, I. Hen, Quantum Sci. Technol. 4 02LT03 (2019)

THE PROBLEM OF IMPLEMENTATION ERRORS

- Annealing schedules that slow down near the minimum gap require precise control
  - Consider the optimal annealing schedule of Roland & Cerf for unstructured search (Grover) (1). The linear annealing schedule is modified such that $s(t/t_f)$ is a non-trivial function
  - Optimal schedule is 'finely' tuned to the location of the minimum gap. If the minimum gap shifts, then the quadratic speedup is lost.

Optimal annealing schedule becomes exponentially slow (with the system size) near the minimum gap at $s = \frac{1}{2}$

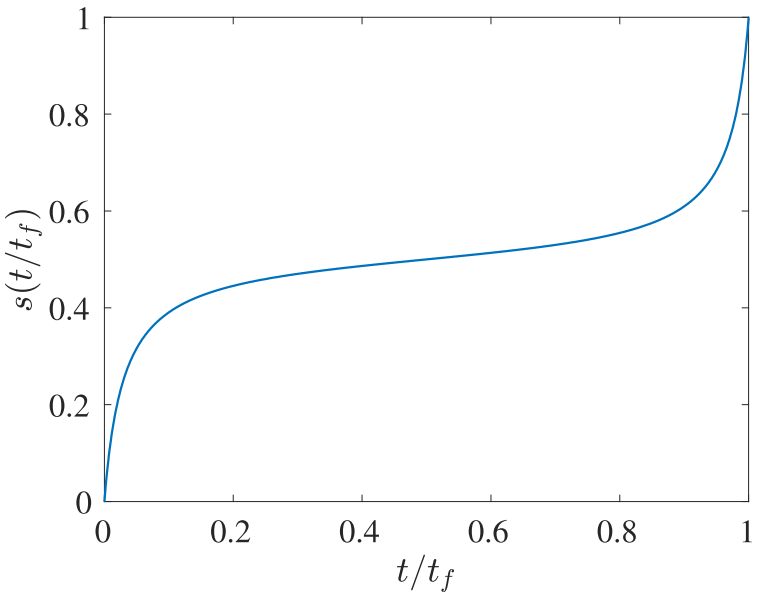

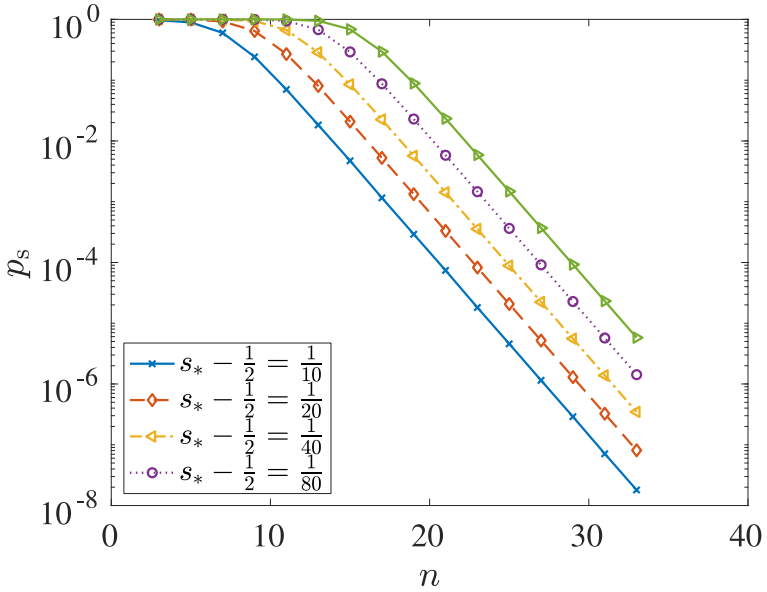

Success of adiabatic evolution decreases exponentially (2) with system size when the minimum gap is shifted away from $s = \frac{1}{2}$ to s_*

(1) J. Roland and N.J.. Cerf, Phys. Rev. A 65, 042308 (2002).
(2) M. Slutskii, TA, L. Barash, I. Hen, arXiv:1904.04420 (2019)

LESSONS FOR THE NISQ ERA

- It is not surprising that in the absence of error correction, performance does not scale.
- Our devices will always have finite size. Errors need only be as small as necessary for the circuits we want to implement.
- Specialized devices to mitigate errors (co-design): algorithms specialized for a given architecture / architectures tailored for a specific algorithm. Reduces overhead of 'embedding' the problem on the device, as well as maximally utilizing 'native' hardware capabilities.
- At small sizes, 'typical' instances may not be representative of 'typical' instances at larger sizes. At small sizes, we should search for the hardest instances to truly understand the limitations of our algorithms and devices.
- Optimism is great, but proper benchmarking against the best classical algorithms is necessary. The goal post is constantly moving, but that's progress!

WHAT EXCITES ME: NON-STOQUASTICITY

- Standard quantum annealing implements a 'stoquastic' Hamiltonian along the entire anneal.
  - Definition: all the off-diagonal elements of the Hamiltonian in the computational basis are ≤ 0. This is basis dependent, so the definition is up to local transformations of the basis.

$$H(s) = -(1-s)\sum_i \sigma_i^x + \lambda\, s(1-s)\sum_{i,j} \sigma_i^x \sigma_j^x + s\, H_p, \qquad s = \frac{t}{t_f} \in [0,1]$$

$$\lambda \leq 0, \text{obviously stoquastic}$$
$$\lambda > 0, \text{possibly non} - \text{stoquastic}$$

- One of the key differences is that the ground state of stoquastic Hamiltonians have only positive amplitudes, whereas the ground states of non-stoquastic Hamiltonians can have complex amplitudes.
- We know that universal adiabatic quantum computing requires non-stoquastic interactions (1).
- Can using such interactions make (non-standard) quantum annealing more powerful?
  - Typically, these interactions make quantum annealing less efficient (2,3)!
  - Only for very rare situations does non-stoquasticity appear to help (4,5). Of the examples we know, the advantage over its stoquastic counterpart is exponential!

(1) J. Biamonte, P. Love, Phys. Rev. A 78, 012352 (2008).
(2) E. Crosson et al., arXiv:1401.7320
(3) L. Hormozi et al, Phys. Rev. B 95, 184416 (2017)
(4) Y. Seki, H. Nishimori, Phys. Rev. E 85, 051112 (2012)
(5) T. Albash, Phys. Rev. A 99, 042334 (2019)

WHAT EXCITES ME: NON-STOQUASTICITY

- Exponential advantage appears to be correlated to the fact that the ground state has more flexibility by allowing for arbitrary amplitudes (1).
- Can we better understand when non-stoquastic interactions can help in quantum adiabatic optimization? Our studies so far show that it is definitely not the typical case.
- Non-stoquastic Hamiltonians are generally very interesting. These are associated with the sign-problem in quantum Monte Carlo. Any progress in this direction will likely have far-reaching implications.
- As part of QEO, we hope to build systems that implement such interactions using superconducting Josephson Phase-Slip Qubits (2). So far these qubits are very sensitive to charge noise, and a lot of progress needs to be made before they are 'device-ready'.

(1) T. Albash, Phys. Rev. A 99, 042334 (2019)
(2) A. J. Kerman, arXiv:1810.01352 (2018)

LOOKING FORWARD BEYOND OPTIMIZATION

- Quantum simulation remains one of the most interesting applications for quantum computing,
- Simulation of non-stoquastic Hamiltonians is highly appealing.
  - Fermionic systems in condensed matter physics, quantum chemistry, high energy physics, …
  - Sign-problem problems: quantum Monte Carlo techniques fail!
  - Other algorithms are agnostic to this: Neural/Tensor network approaches. Suggests that not all non-stoquastic systems are hard to simulate. When they do get hard (when long-range entanglement is present) is also when we might expect our quantum device to be challenged.

Suggests there is an open question to be addressed about when we can meaningfully (in the absence of fault tolerant error correction) beat classical simulations of such systems.
Proper benchmarking needs to be done to make meaningful comparisons.

SOME SPECULATION

- So far, experience has taught us that whenever things get interesting from a quantum point of view, things also get challenging experimentally (1). Is there something fundamental about robustness/sensitivity and a demonstration of a quantum advantage?

	Classically Easy	Classically Hard
Experimentally Quantum Robust	Highly likely?	How likely is this?
Experimentally Quantum Sensitive	Useful for verification?	Highly likely?

Simulating non-stoquasticity systems seems to be the ideal playground to explore such questions.

(1) P. Hauke, et al., Rep. Prog. Phys. 75, 082401 (2012)

RESEARCH NEEDS

- Benchmarking quantum simulation will require a broad range of simulation and algorithmic techniques. We could try to build these in-house, but this is an excellent (and easy) opportunity for collaboration.
- Connecting to applications of interest to the 'real world' requires domain expertise. Inter-disciplinary should be more than just a buzz word.
- Understanding the quantum information processor is key to understanding and improving its performance (especially for error mitigation/suppression before full quantum error correction). Opportunities for lots of open-quantum system modeling and interactions with experimentalists.

A great example combining all these features:
IBM's recent-ish VQE paper (Nature 549, 242 (2017))
Used domain expertise to map relevant degrees of freedom of H_2. LiH, BeH_2 to 5 qubits, open quantum system simulations give excellent agreement with device output, showing how decoherence is a significant limiter of performance