Challenges in Quantum Computation, UC Berkeley June 2018

Challenges in Adiabatic Optimization

Elizabeth Crosson

California Institute of Technology

Caltech

Transverse-field adiabatic optimization

(Farhi, Goldstone, Gutmann, Sipser. 2000)

- Minimize a cost function $f : \{0,1\}^n \to \mathbb{R}$ by sampling the ground state of an n-qubit Hamiltonian,

$$H_p = \sum_{z \in \{0,1\}^n} f(z)|z\rangle\langle z|$$

- Initialize the qubits in the ground state of a uniform transverse field $H_B = -\sum_{i=1}^{n} X_i$ and interpolate from H_B to H_p,

$$H(s) = (1-s)H_B + s\,H_p \quad , \quad 0 \leq s \leq 1$$

- **Adiabatic theorem:** running for time $\mathrm{poly}(n, \Delta^{-1})$, where $\Delta = \min_s E_1(s) - E_0(s)$ is the minimum spectral gap of $H(s)$, suffices to prepare the ground state of H_p.

TF-AO $\subseteq$ Adiabatic Optimization $\subseteq$ Adiabatic Computation

- Adiabatic optimization can also use different paths e.g.

$$H(s) = (1-s)H_B + s(1-s)H_E + s\,H_p \quad , \quad 0 \leq s \leq 1.$$

- For some example cost functions, a nontrivial H_E can improve the min gap Δ from $\mathcal{O}(2^{-n})$ to $\Omega(1)$. (FGG, 2002).

- **Adiabatic computation:** Using more general local Hamiltonian paths $H(s)$, ground state adiabatic evolution + measurement + classical post-processing is a universal model of quantum computation. (Aharanov et al., 2007).

- In universal AQC constructions the final Hamiltonian $H(s=1)$ is a Feynman-Kitaev circuit Hamiltonian.

Heuristic reasons for interest in QAO

- Implementable as an analog algorithm, with expectations of inherent robustness to errors. (Childs, Farhi, Preskill, 2004)
- Optimization principles are common in nature, and have inspired valuable classical optimization algorithms (e.g. Simulated Annealing, Kirkpatrick et al., 1983. +30k citations)
- Intuition for quantum tunneling to speed up exploration of rugged energy landscapes in the classical cost function.
- Tunneling intuition led to toy problems in which adiabatic optimization is exponentially faster than classical SA (FGG '02) or any local search algorithm (Reichardt '04).

Exponential speedup over classical SA

- **Spike cost function:** bit-symmetric cost function with a large energy barrier that creates a local minimum.

$$f(w) = \begin{cases} |w| + n^a & n/4 - n^b/2 \leq |w| \leq n/4 + n^b/2 \\ |w| & o.w. \end{cases}$$

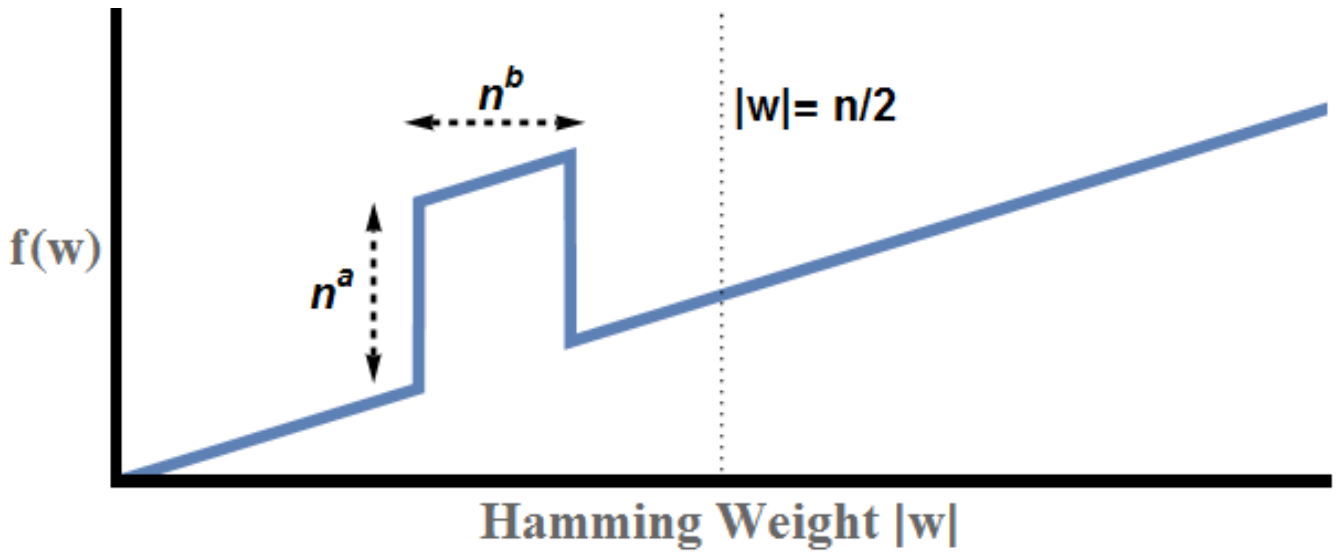

- Takes time $\Omega(2^{n^a})$ to solve with local search algorithms, but QA takes $O(n)$ time when $a + b < 1/2$ (Reichardt '04).

- **Bad news:** No evidence of speedup over best-known classical optimization algorithms using QAO with local Hamiltonians.
- Grover speedup can be obtained using oracle Hamiltonians resembling QAO. (Roland and Cerf, 2001).
- Glued-trees speedup can be obtained using adjacency matrix Hamiltonian oracle and a nearly-degenerate ground-space. (Nagaj, Somma, Keiferova, 2012)
- QAOA superficially resembles QAO and briefly achieved the best-known approximation ratio guarantee for MAX-E3-LIN2.
- Inspiring rigorous algorithms: reverse annealing (Smelyanskiy et al. 2018), short-path optimization (Hastings, 2018)

Why is there no speedup? Is it because H is stoquastic?

- Even empirical benchmarks of TF-AO do not indicate any speedup over the best classical algorithms.
- Most commonly cited reason for the lack of a speedup is that TF-AO Hamiltonian is **stoquastic** in the computational basis.
- Hamiltonian $H = \sum_i H_i$ is stoquastic if in some local basis $\mathcal{B}$ the terms H_i all have matrix entries that are zero or negative,

$$\langle x|H|y\rangle \leq 0 \quad , \quad \forall x, y \in \mathcal{B} \text{ with } x \neq y.$$

- Stoquastic = "quantum" + "stochastic"
- H "doesn't have a sign problem", and its equilibrium states are on the border of quantum and classical physics.

Transverse Ising Models

- **Transverse Ising models** are stoquastic in the Z basis,

$$H = -\sum_{i=1}^{n} X_i - \sum_{i,j} Z_i Z_j \quad , \quad X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad , \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

- So are generalized TIM with disordered interactions,

$$H = -\sum_{i} \Gamma_i X_i + \sum_{i,j} \alpha_{ij} Z_i Z_j + \sum_{i} b_i Z_i \quad , \quad \Gamma_i > 0$$

- Theorem: generalized TIM are universal for stoquastic adiabatic computation. Proof uses perturbative gadgets. (Bravyi, Hastings '14, Cubitt, Montanaro, Piddock '16).

Perron Frobenius theorem: amplitudes in equilibrium

- If H has all real and non-positive matrix entries, then $A = -\beta H$ is a non-negative matrix. Expand e^A as a series,

$$e^A = 1 + A + \frac{A^2}{2} + ...$$

- Every term in the series is a nonnegative matrix, therefore $e^{-\beta H}$ is a nonnegative matrix.

- Since $\lim_{\beta\to\infty} e^{-\beta H} = |\psi\rangle\langle\psi|$, there is a choice of global phase which gives the ground state nonnegative amplitudes.

- If H is an irreducible matrix, then all of the ground state amplitudes are positive, $\psi(x) > 0$ for all $x \in \mathcal{B}$. (no nodes)

Euclidean path integrals

- We can expand the partition function as a "path integral",

$$Z = \mathrm{tr}\left(e^{-\beta H}\right) = \sum_{x \in \mathcal{B}} \langle x| \left(e^{-\frac{\beta H}{L}}\right)^L |x\rangle = \sum_{x_1,\ldots,x_L \in \mathcal{B}} \prod_{i=1}^{L} \langle x_i| e^{-\frac{\beta H}{L}} |x_{i+1}\rangle$$

- Since the "propagator" $e^{-\frac{\beta H}{L}}$ is a nonnegative matrix, every path $(x_1, ..., x_L)$ contributes a positive weight to this sum.

- Define a probability distribution on the space of paths,

$$\pi(x_1, ..., x_L) = \frac{1}{Z} \prod_{i=1}^{L} \langle x_i| e^{-\frac{\beta H}{L}} |x_{i+1}\rangle$$

- For large β, this π has the ground state distribution as a marginal, $\psi(x)^2 = \sum_{x_2,\ldots,x_L} \pi(x, x_2, ..., x_L)$.

Quantum Monte Carlo

- **Idea:** MCMC for stoquastic path integrals.
- Ground state probability distribution of H becomes a marginal of a Gibbs distribution in a larger state space,

$$\pi(x_1, ..., x_L) = \frac{e^{-E(x_1,...,x_L)}}{Z} \quad , \quad x_i \in \{0,1\}^n$$

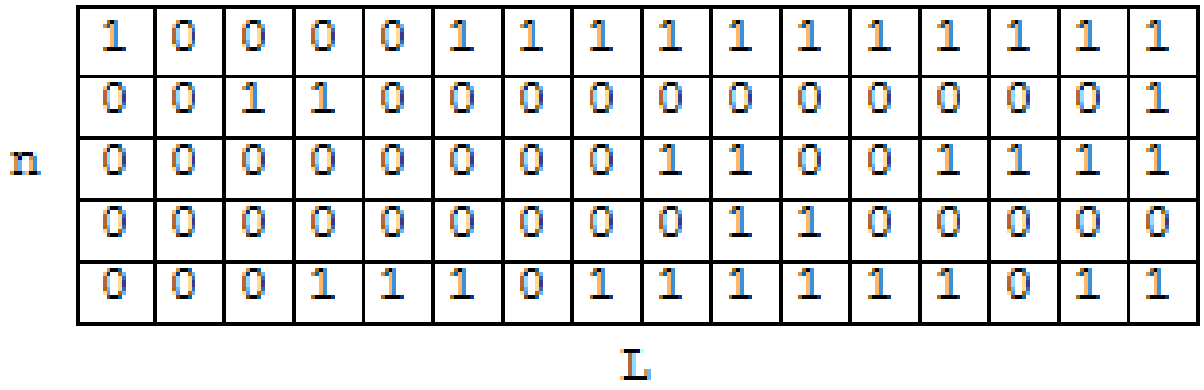

- The path integral is a faithful approximation, as long as we approximately sample from π.
- QMC works works well in practice, but convergence of the MCMC is not theoretically explained in general.

D-Wave: the original NISQ device

- D-Wave implements a noisy version of transverse-field adiabatic optimization called quantum annealing.
- Single qubit coherence times $\sim$1ns , annealing times $\sim 1\mu$s, substantial open system effects, calibration errors
- Successfully solves 2000 bit optimization instances, works better than theory might predict (confirming some sense of inherent robustness).
- For instances designed to give D-Wave every advantage over its classical competitors, it matches the performace of the best implementations of the best algorithms on a modern single-core CPU.

Empirical Benchmarks

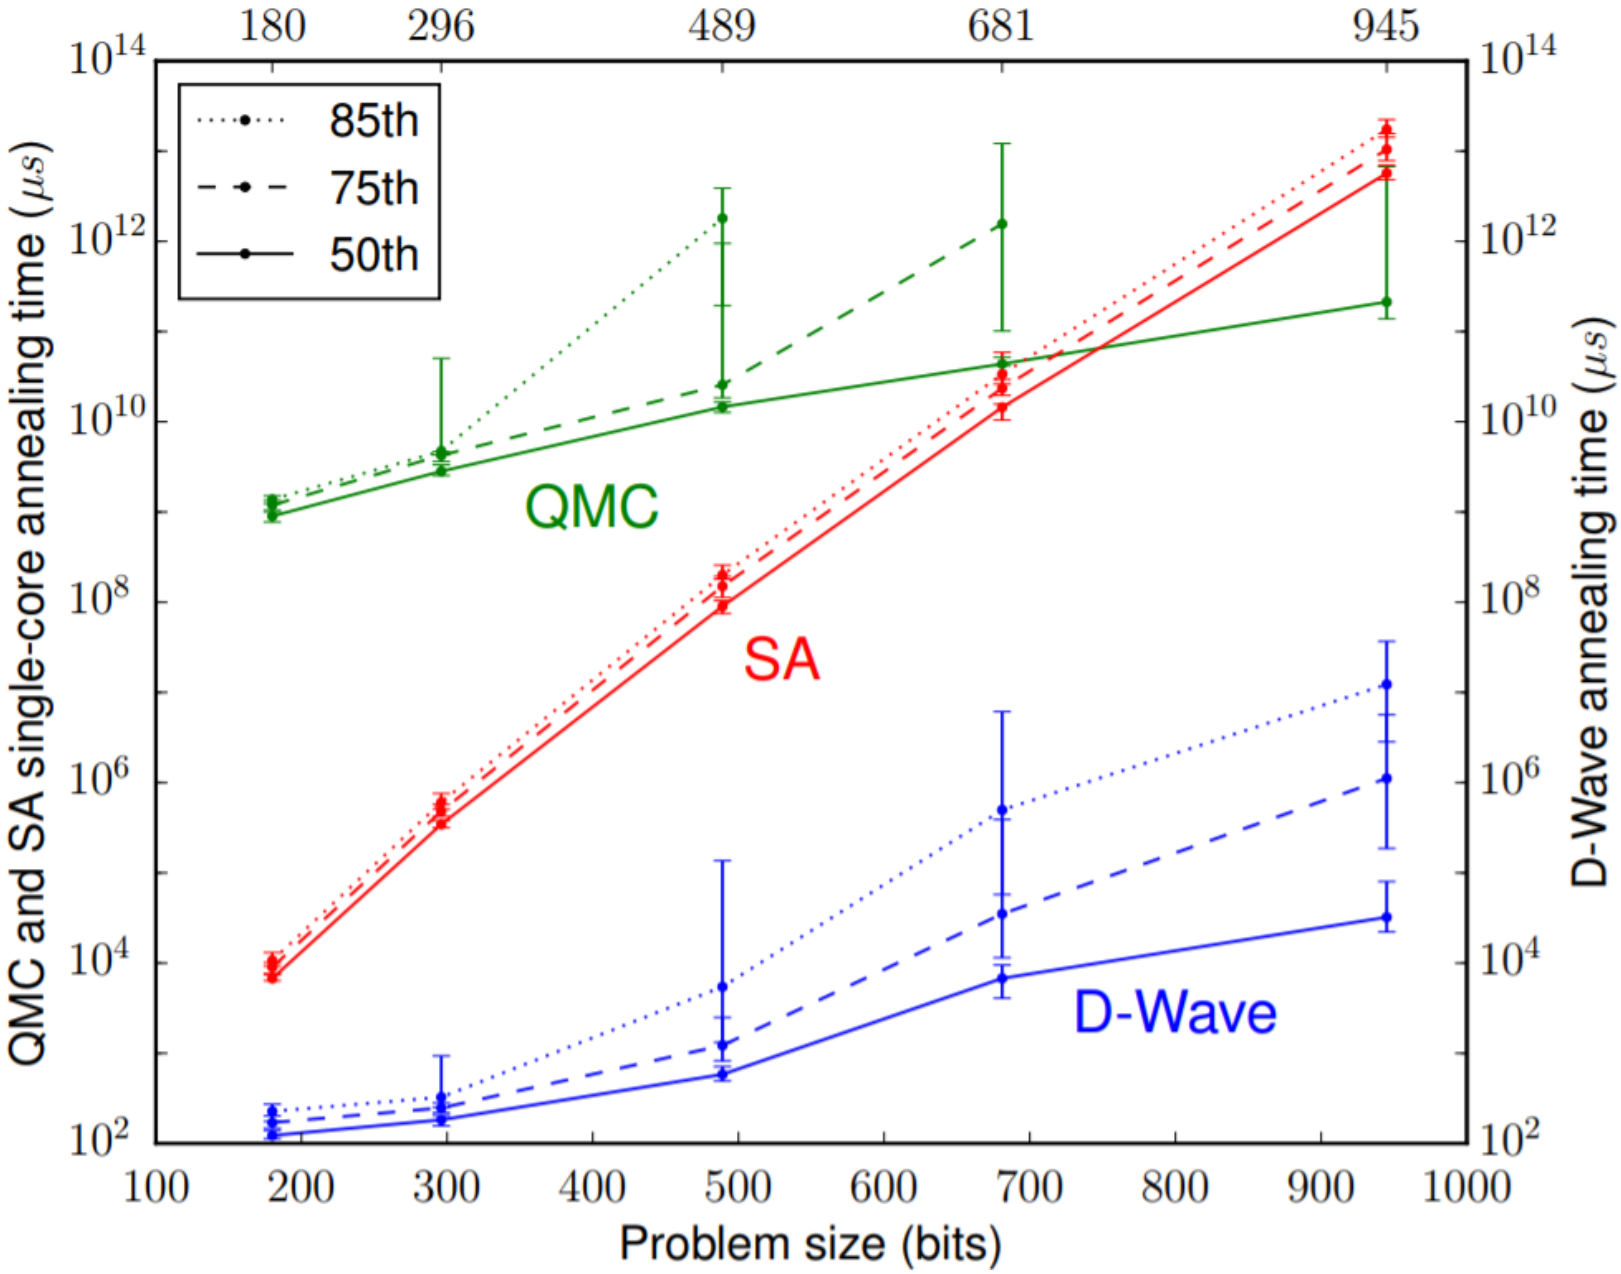

The stoquastic simulation conjecture

- **Conjecture:** stoquastic adiabatic computation can be classically simulated in time $\mathrm{poly}(n, \Delta^{-1})$, in the sense of sampling the ground state probability distribution (GSPD) in the computational basis.

- Crucial that $\Delta := \min_s \Delta(s)$ for a sensible conjecture, since NP-hard instances with a unique ground state retain a constant gap when perturbed by a transverse field.

- Markov chains are designed to forget where they came from, so making essential use of the adiabatic path in a simulation attempt is quite difficult.

- It turns out that "warm starts" are not really enough, because $\langle x|\psi\rangle$ is $\Omega(n^{-n})$ for generalized TIM ground states $|\psi\rangle$ and all $x \in \{0,1\}^n$.

Status of the stoquastic simulation conjecture

- Most general result to date is poly-time simulation of frustration-free stoquastic AQC (Bravyi, Terhal, 2008).

- All other rigorous progress is on "easy" cases: ferromagnetic models on any graph and any temperature (Bravyi and Gosset), spike Hamiltonians and 1D models at $\beta = \mathcal{O}(\log n)$ temperature (EC and Harrow).

- All existing QMC algorithms are inadequate for the most general form of the conjecture, due to topological obstructions and L_1 vs L_2 obstructions. (Hastings 2013).

Quantum supremacy with stoquastic ground states

- Stoquastic adiabatic computation cannot be universal (unless PH collapses to 3rd level) because approximating stoquastic path integrals is in PostBPP. (Bravyi et al. 2005)

- However, Bravyi et al. define H to be stoquastic if there is *any* choice of local basis in which the condition on the matrix entries of H is satisfied.

- Therefore theoretical QS could potentially be obtained by sampling a stoquastic Hamiltonian in a rotated basis.

- Indeed, output of an IQP circuit is the ground state of local Hamiltonian $UHU^\dagger$ where $H = -\sum_i X_i$ and U is depth 3.

Quantum supremacy with stoquastic ground states

- More directly, 1-local rotations suffice to make the 2D cluster state Hamiltonian used in MBQC explicitly stoquastic.

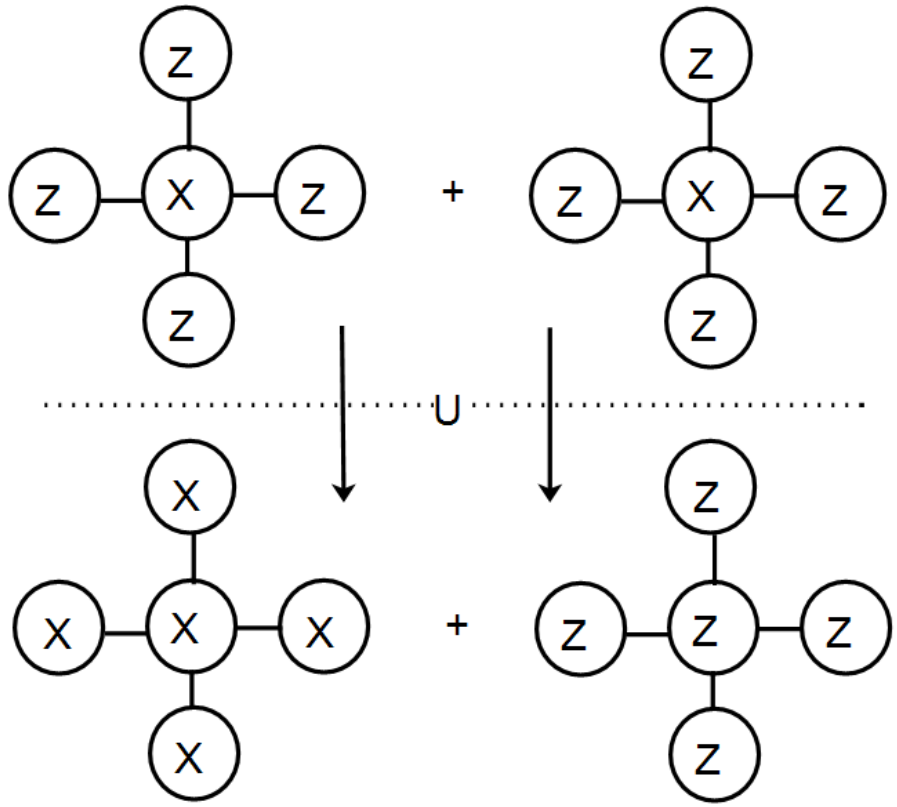

- Similar observation in "Quantum speedup in stoquastic adiabatic quantum computation" (Fujii 2018).

How can we evaluate nonstoquastic QAO?

- **Isoperimetric inequality:** relates the geometry of the ground state probability distribution (GSPD) to the spectral gap.

- Generalizes a known result for Markov chains and stoquastic Hamiltonians to **nonstoquastic** Hamiltonians.

- Furthers our understanding of the probability distributions that arise from nonstoquastic ground states.

- *Quantum ground state isoperimetric inequalities for the energy spectrum of local Hamiltonians*.
(EC & J. Bowen, arXiv:1703.10133).

Ground states and weighted graphs

- Think of the labels of a basis $\mathcal{B}$ as vertices of a graph, and the ground state $|\psi\rangle$ as a probability distribution on $\mathcal{B}$,

$$x \in \mathcal{B} \quad \Rightarrow \quad \pi(x) = |\psi(x)|^2.$$

- Connect two vertices $x, y \in \mathcal{B}$ by an *unweighted* edge if the corresponding Hamiltonian matrix entry is nonzero,

$$x, y \text{ connected by an edge} \quad \Leftrightarrow \quad \langle x|H|y\rangle \neq 0.$$

- Define the interior boundary of a set of vertices $S \subseteq \mathcal{B}$ as the vertices in S connected to vertices outside of S,

$$\partial S = \{x \in S : \exists y \notin S \text{ with } \langle x|H|y\rangle \neq 0\}.$$

Isoperimetric Inequality for Quantum Ground States

- **Notation:** let H have ground state energy $E \geq 0$ and operator norm $\|H\|$, and let $\Delta_H := E_1 - E$ be the spectral gap of H.

- **Theorem:** if $\mathcal{B}$ is an arbitrary basis and π is the GSPD of H in the basis $\mathcal{B}$, then any subset $S \subset \mathcal{B}$ with $\pi(S) \leq 1/2$ satisfies

$$\Delta_H \leq 2(\|H\| - E)\frac{\pi(\partial S)}{\pi(S)}.$$

- Depends on the range of k-local terms in H (through ∂S), but not on the details of the Hamiltonian couplings.

- For a given π the inequality constrains the spectral gap of any k-local H having π as its GSPD in any choice of local basis.

Conclusion

- Simulated quantum annealing works well for stoquastic QAO, so nonstoquastic H seem necessary for quantum speedup.
- Sampling a stoquastic H in a locally rotated basis can yield theoretical and practical quantum supremacy.
- Adiabatic computation with nonstoquastic H is universal, but this doesn't mean that nonstoquastic H will improve the performance of adiabatic optimization.
- Some explicit distributions inevitably take a long time to precisely sample using purely adiabatic evolution $\Longrightarrow$ spectral gap condition is too pessimistic, leaving the ground state is advantageous (and so are measurements and ancillas).
- *Thank you for your attention!*