(sub)Exponential quantum speedup for the Guided Stoquastic Hamiltonian problem

Yassine Hamoudi

Ongoing work with Shrinidhi Teganahally Sridhara and Yvan Le Borgne

Hamiltonian complexity

How hard is it to compute ground state information
of physically relevant quantum systems?

Hardness can be quantified in terms:

- Complexity classes (BPP, BQP, QMA…)
- Resource lower bounds (space, queries, gates…)
- Structural properties of the ground state (entanglement/correlation measures, symmetries…)

This talk:

∃ *A "natural" Hamiltonian problem that is much harder for classical computers (than for quantum computers), even when assisted with some "basic" quantum resources*

The Ground State Computation problem

Given a Hamiltonian H, compute the eigenvector $|\psi_H\rangle$ ("ground state") corresponding to the smallest eigenvalue λ_H ("ground energy")

Closely related to Ground Energy estimation: *compute* $\lambda_H \pm \epsilon$

QMA-hard in general (⇒ no efficient quantum algorithm is believed to exist)

Becomes much easier when the following auxiliary resource is available:

Guiding state: State $|\tilde{\psi}\rangle$ that has large overlap with $|\psi_H\rangle$
(a.k.a. warm start)

The Ground State Computation problem

Given a Hamiltonian H, compute the eigenvector $|\psi_H\rangle$ ("ground state") corresponding to the smallest eigenvalue λ_H ("ground energy")

A simple quantum algorithm:

Run Quantum Phase Estimation on the guiding state $|\tilde{\psi}\rangle$ with enough precision to distinguish component $|\psi_H\rangle$ from the others, and collapse onto it by measuring the phase estimate

$$|\tilde{\psi}\rangle \xmapsto{\text{QPE}} \langle\psi_H|\tilde{\psi}\rangle|\psi_H\rangle|\lambda_H \pm \epsilon\rangle + \sqrt{1-\langle\psi_H|\tilde{\psi}\rangle^2}\,|\psi_H^\perp\rangle|\ldots\rangle \xmapsto{\text{Measure}} |\psi_H\rangle$$

The Ground State Computation problem

Given a Hamiltonian $H \in \mathbb{C}^{2^n \times 2^n}$, compute the eigenvector $|\psi_H\rangle$ ("ground state") corresponding to the smallest eigenvalue λ_H ("ground energy")

Run Quantum Phase Estimation on the guiding state $|\tilde{\psi}\rangle$ with enough precision to distinguish component $|\psi_H\rangle$ from the others, and collapse onto it by measuring the phase estimate

Under the working assumptions:

- The Hamiltonian is efficiently simulatable — *local, sparse, … and* $\|H\| \leq \mathrm{poly}(n)$
- The Hamiltonian is gapped — *difference* $\geq 1/\mathrm{poly}(n)$ *between* λ_H *and 2nd eigenvalue*
- The guiding state has good overlap with $|\psi_H\rangle$ — $|\langle\psi_H|\tilde{\psi}\rangle| \geq 1/\mathrm{poly}(n)$

… the quantum algorithm runs in time $\mathrm{poly}(n)$

Can we reasonably expect a guiding state to be known?

Lots of quantum algorithms/heuristics for preparing possibly-good guiding states:

Variational algorithms, quantum state ansätze (Matrix Product States, QAOA ansatz, …), Simulated Annealing, Quantum Adiabatic Algorithm…

<u>Example:</u> bootstrapping a guiding state

If we slightly perturb a gapped Hamiltonian, its ground state should remain approximately the same

$$H_0 \longrightarrow H_0 + \delta H_1 \longrightarrow H_0 + \delta H_1 + \delta H_2 \dashrightarrow H$$

Ground state $|\psi_0\rangle$ — Ground state $|\psi_1\rangle \approx |\psi_0\rangle$ — Ground state $|\psi_H\rangle$

"Easy" Hamiltonian — "Target" Hamiltonian

Can we reasonably expect a guiding state to be known?

Example (continued): bootstrapping a guiding state via the Adiabatic Algorithm

H_0: starting Hamiltonian with easy-to-prepare ground state (ex: transverse field Hamiltonian)

$$H_0 = -\sum_i X_i$$

H: target Hamiltonian whose ground state is sought (ex: QUBO Hamiltonian)

$$H = -\sum_{i,j} (Z_i - Z_j)^2$$

$H(t)$: interpolating Hamiltonian with $H(0) = H_0$ and $H(1) = H$ (ex: line interpolation)

$$H(t) = (1-t)\,H_0 + t\,H$$

Can we reasonably expect a guiding state to be known?

Example (continued): bootstrapping a guiding state via the Adiabatic Algorithm

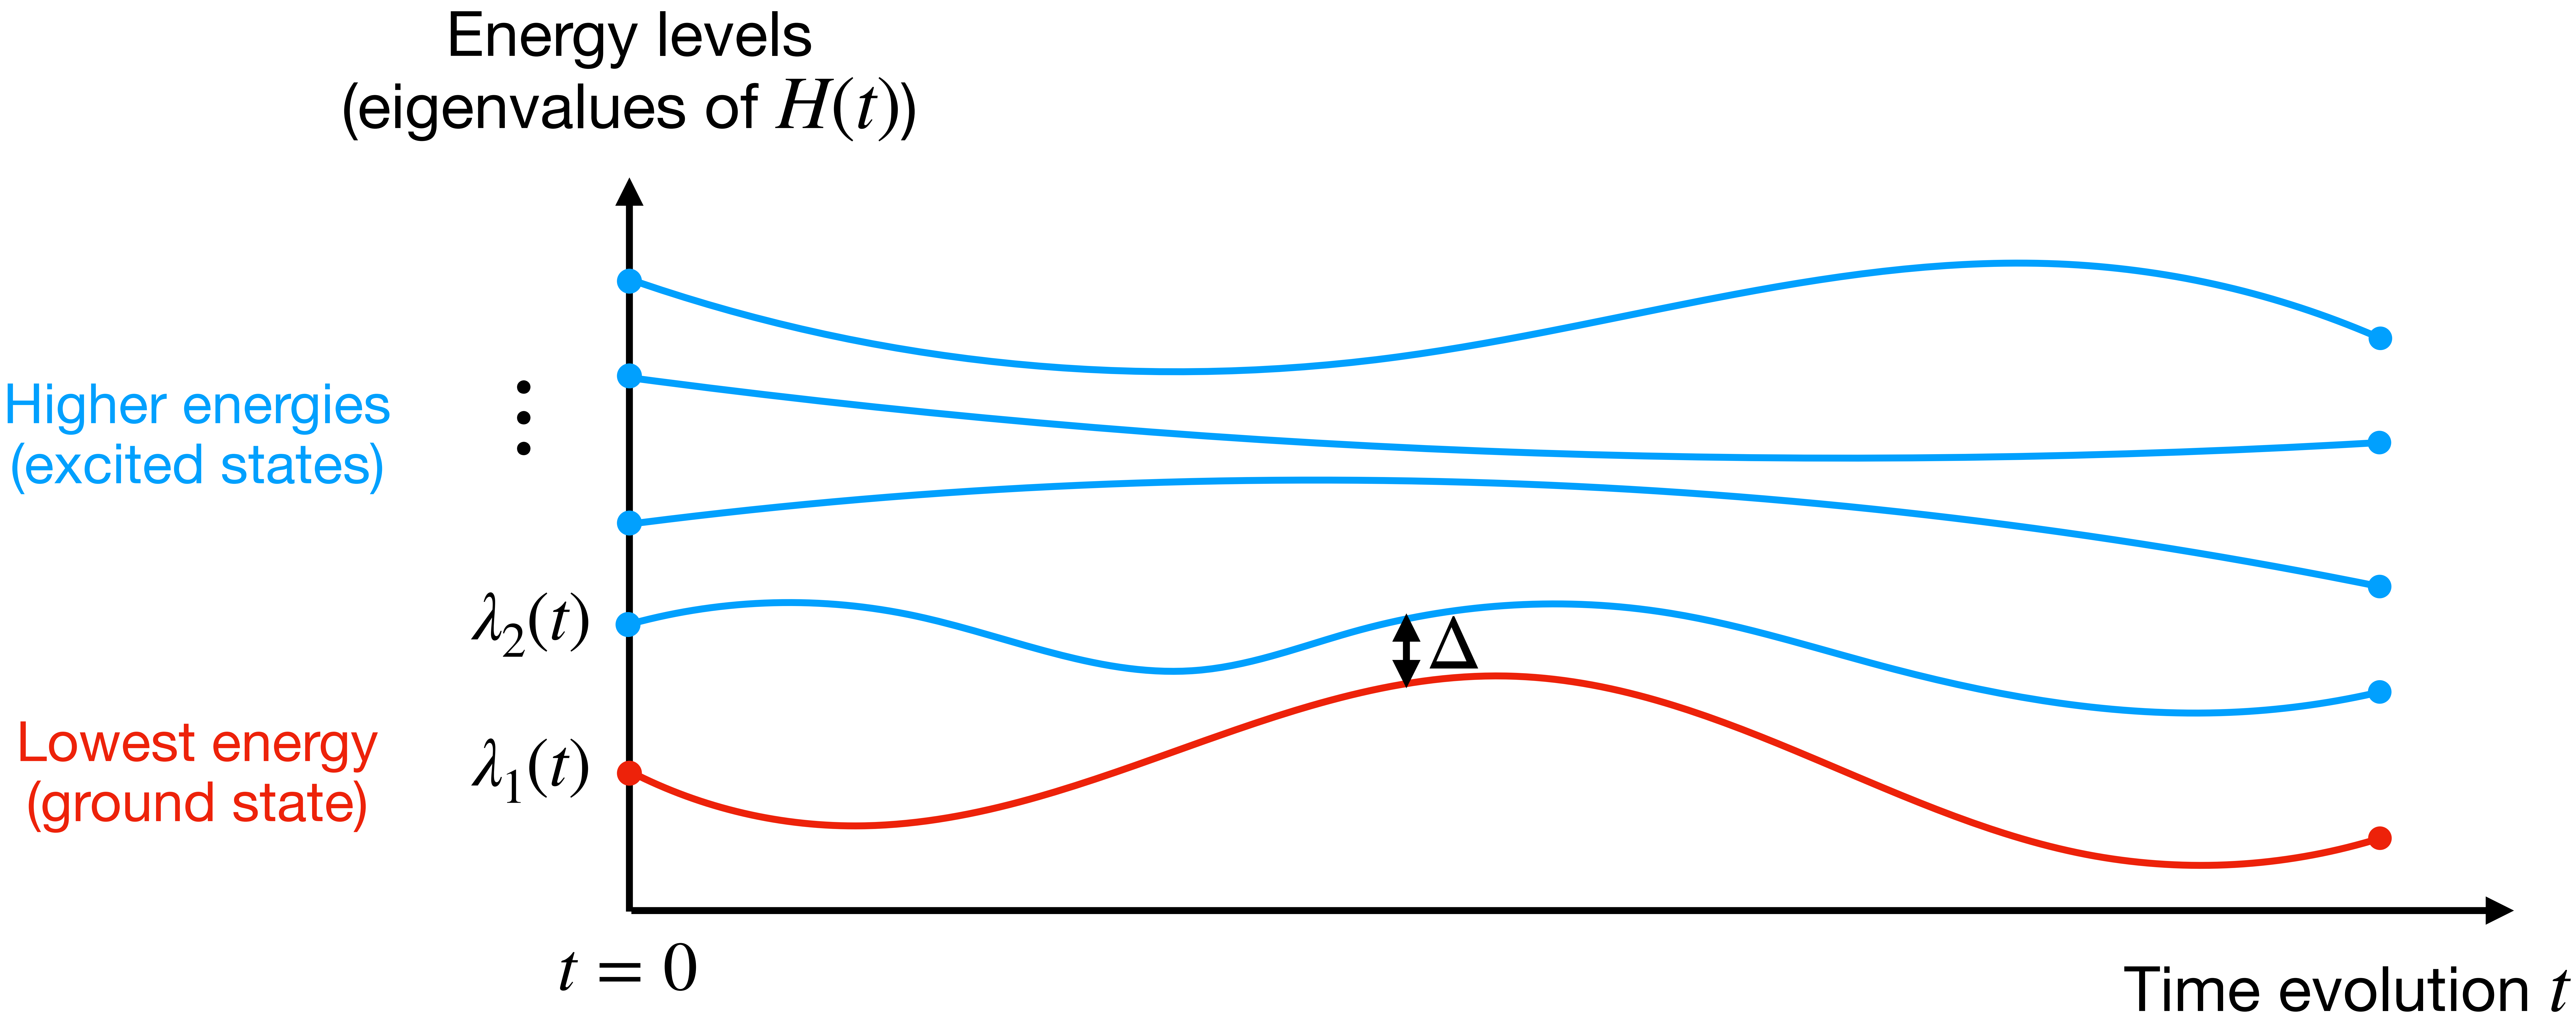

Can we reasonably expect a guiding state to be known?

Example (continued): bootstrapping a guiding state via the Adiabatic Algorithm

If a system is initialized in the ground state of a Hamiltonian $H(t)$ that evolves slowly over time, then it remains in the instantaneous ground state.

$$i\frac{d|\psi(t)\rangle}{dt} = H(vt)|\psi(t)\rangle$$

evolution slowed-down at speed $v \in (0,1]$

Under sufficient regularity assumptions on $H(t)$, the cost is proportional to $1/v \approx 1/\Delta^2$ where:

$$\Delta = \min_{0 \leq t \leq 1} \lambda_2(t) - \lambda_1(t)$$

(minimum spectral gap)

The Ground State Computation problem

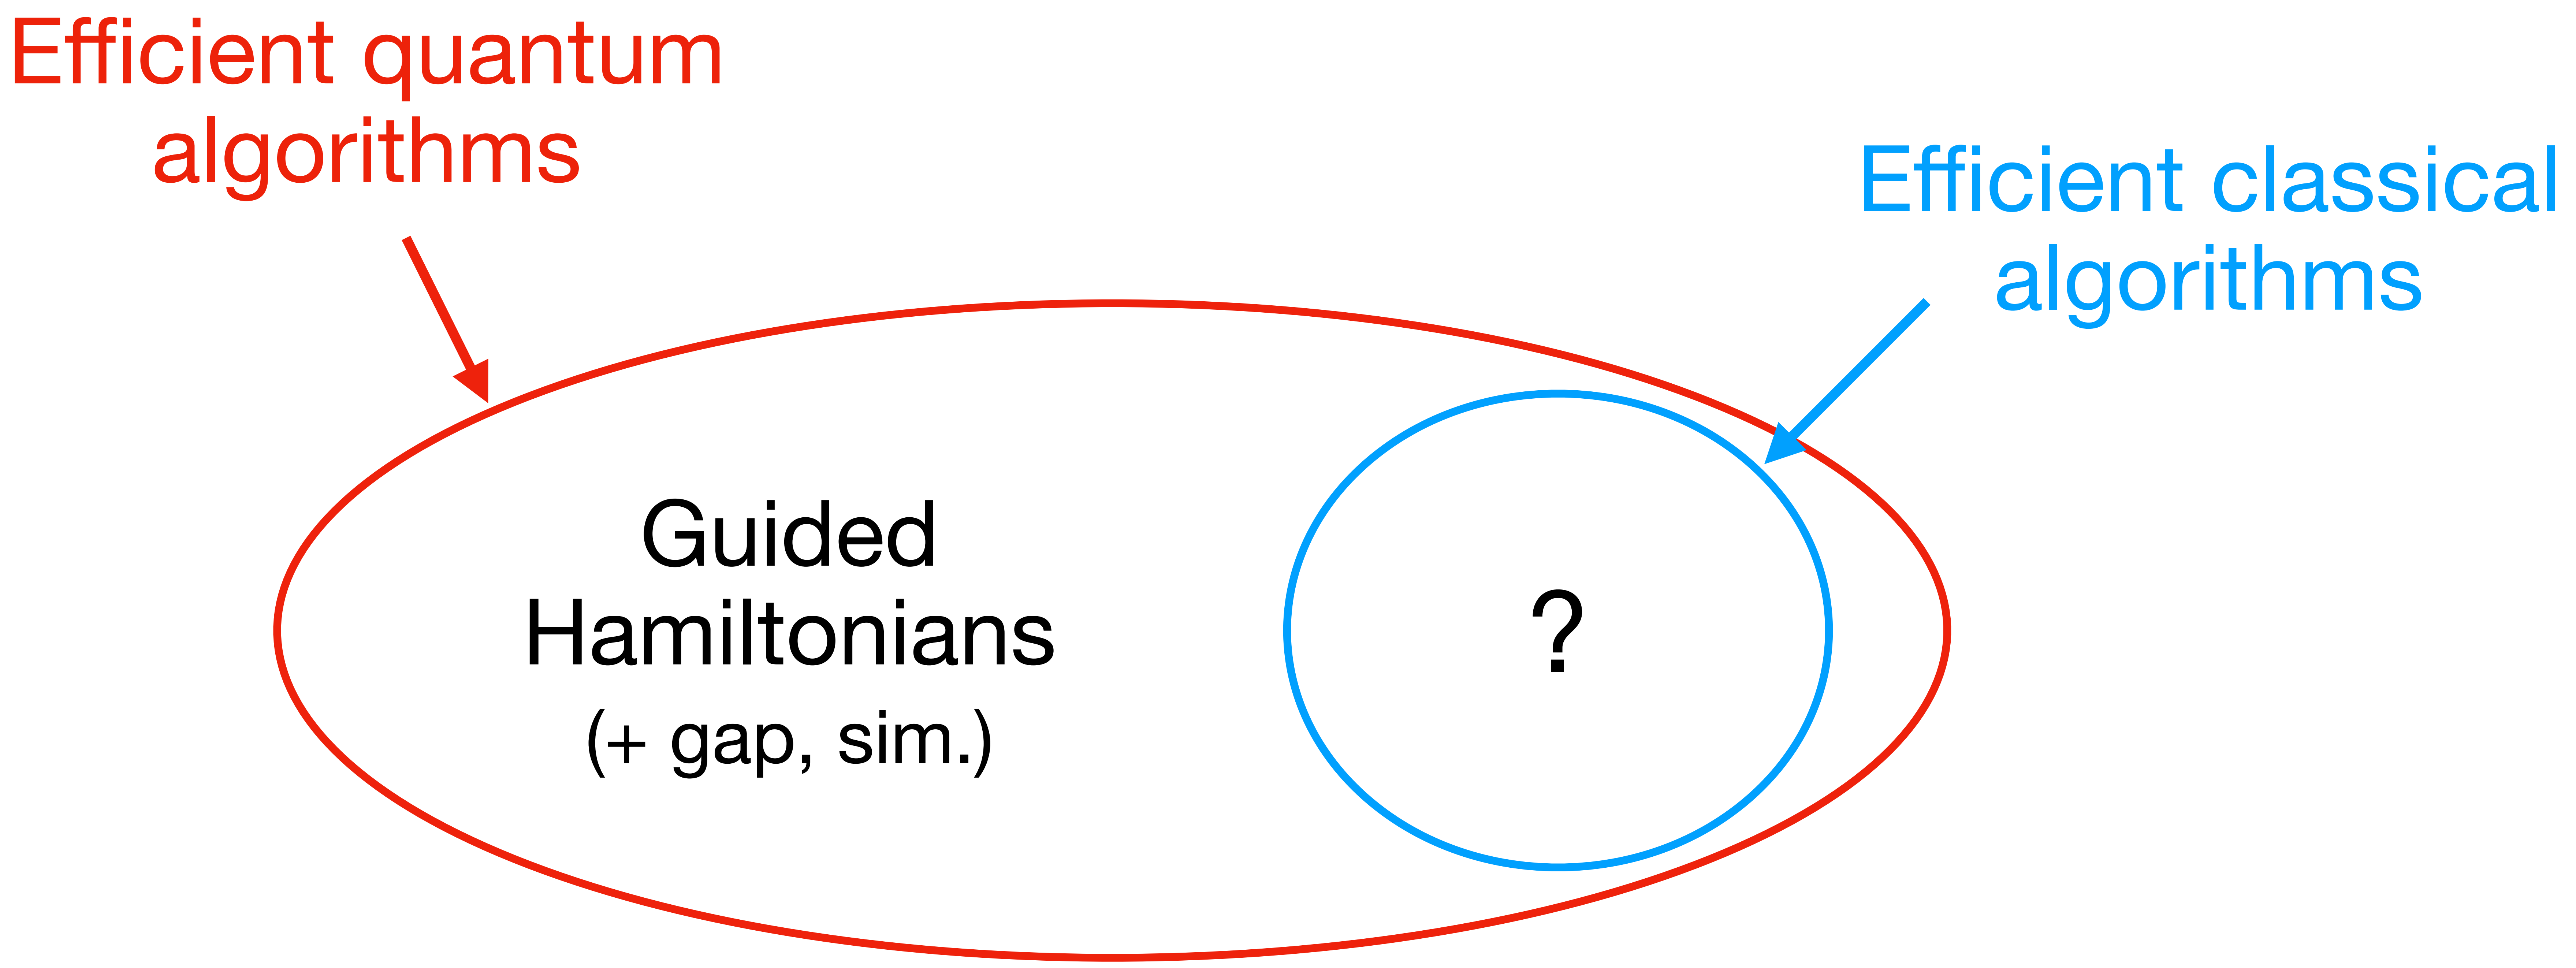

- For classical algorithms, what should it mean to “be given a guiding state” and to “prepare the ground state”?

- Does the problem remain hard for classical algorithms when they are guided?

Classically-Guided Hamiltonians

Quantum algorithms rotate the guiding state toward the ground state

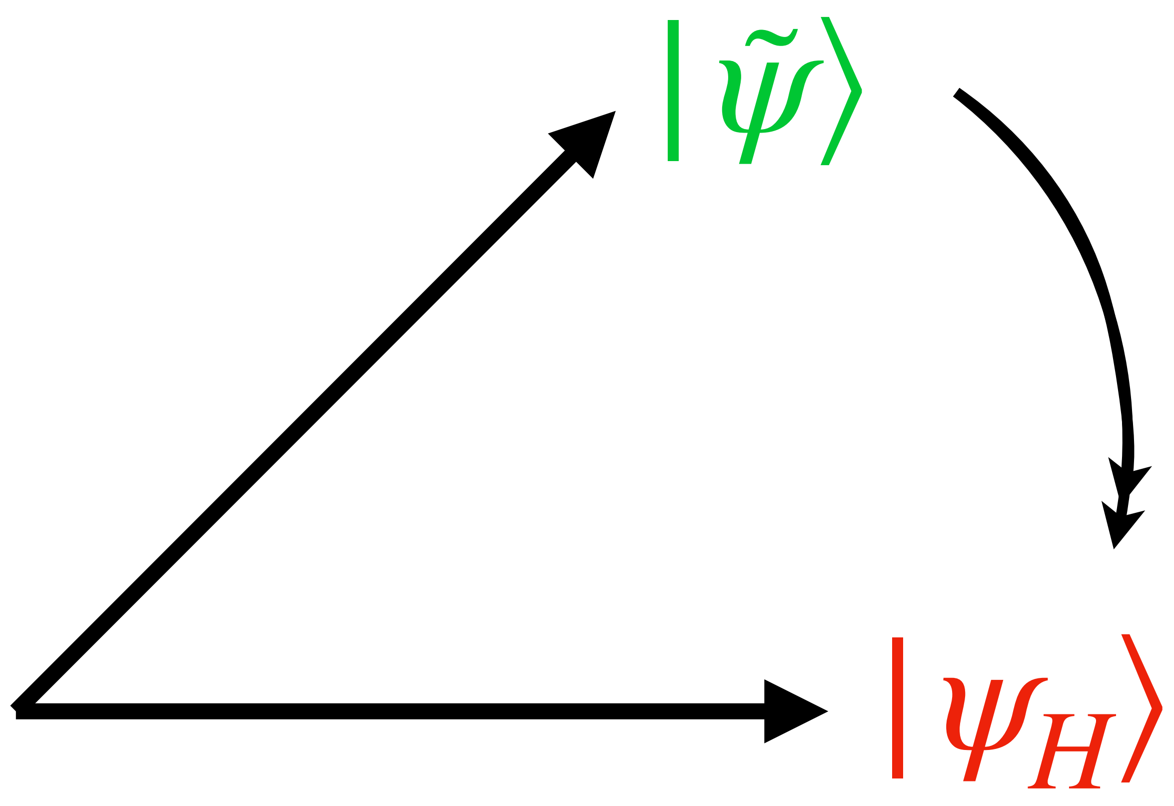

Classical algorithms are asked to realign the computational-basis distributions

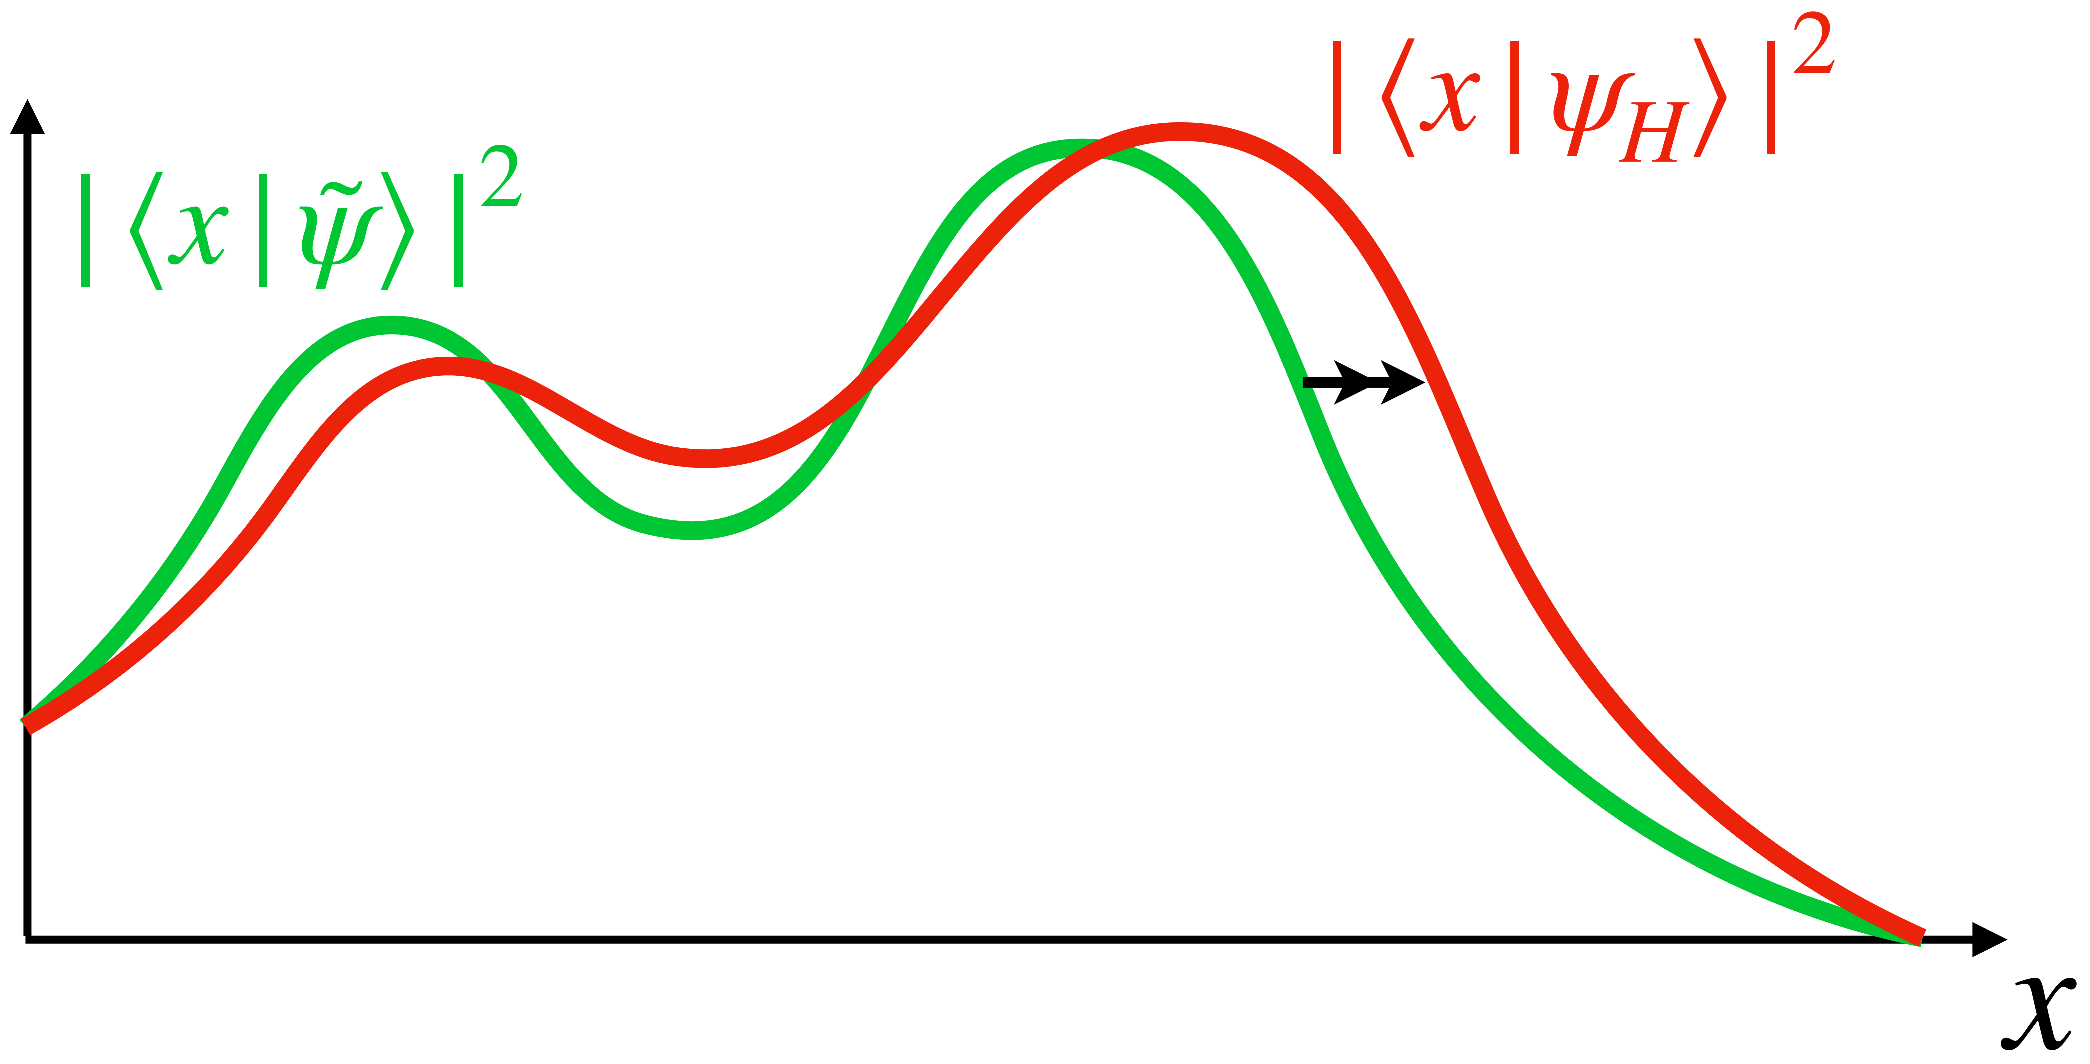

Classically-Guided Hamiltonians

Fix a guiding state $|\tilde{\psi}\rangle$ *and an integer* $T \leq \mathrm{poly}(n)$

<table>
<tr><th>Problem</th><th>Input</th><th>Output</th></tr>
<tr><td>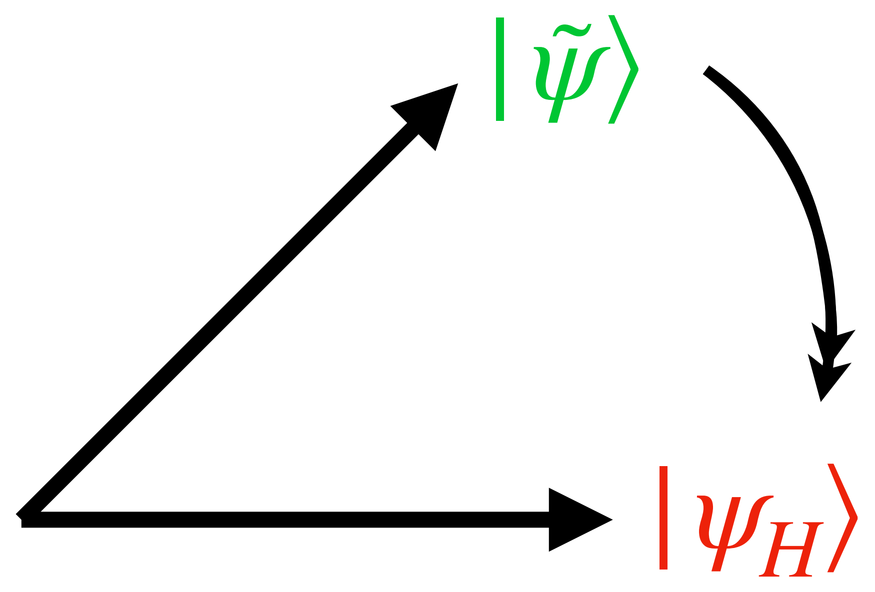
</td><td>- Description of H
- Few copies of the guiding state, $|\tilde{\psi}\rangle^{\otimes T}$</td><td>One copy of the ground state $|\psi_H\rangle$</td></tr>
<tr><td>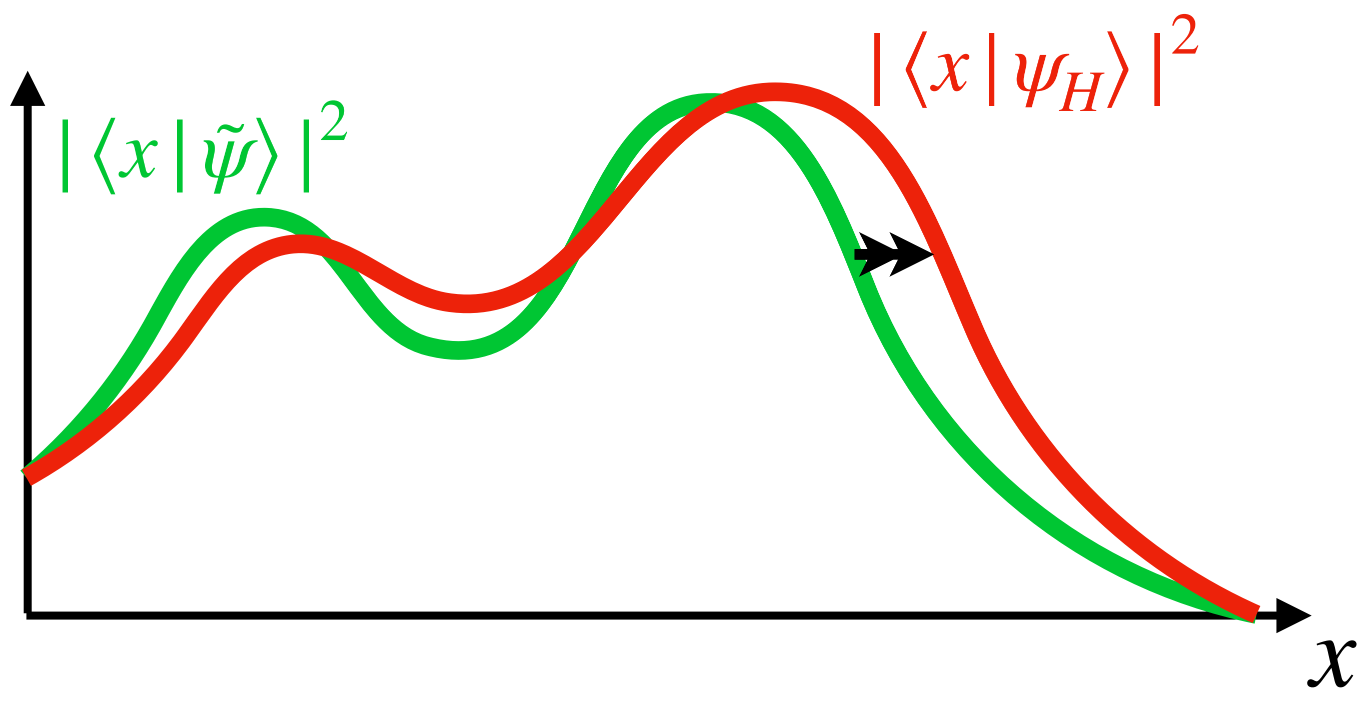
</td><td>- Description of H
- Few i.i.d. samples drawn as $x_1, x_2, \ldots, x_T \sim (|\langle x|\tilde{\psi}\rangle|^2)_x$</td><td>One sample drawn as $x_H \sim (|\langle x|\psi_H\rangle|^2)_x$ independent from $x_1, x_2, \ldots, x_T$

(relax.: at distance ϵ from this distribution)</td></tr>
</table>

Classically-Guided Hamiltonians

Potential applications: improved dequantized algorithms

Dequantized algorithms began with the work of [Tang '19], which showed that some believed quantum speedups were largely achievable by classical algorithms given appropriate access (essentially, L_2 samples) to the input.

Classical Algorithms for Constant Approximation of the Ground State Energy of Local Hamiltonians

Author François Le Gall

Ground energy est. $\lambda_H \pm \epsilon$ *with classical complexity* $\mathrm{poly}\left(1/|\langle\psi_H|\tilde{\psi}\rangle|^{1/\epsilon}, n\right)$ *for guided local Hamiltonians*

⇒ "Boosting" the overlap of the guiding state would improve the running time

Another example:
A dequantized adiabatic algorithm?

A testbed: Stoquastic Hamiltonians

An Hamiltonian is stoquastic (a.k.a. "sign-free") if its off-diagonal terms are non-positive

- Ground state has non-negative real amplitudes → display less interference effects
- Adiabatic computation is often performed using stoquastic Hamiltonians

Example: $H(t) = -(1-t)\sum_i X_i + t \sum_{x\in\{0,1\}^n} f(x)|x\rangle\langle x|$

↑ Classical function to optimize (ex: MAX-SAT)

- Classical simulation methods tend to perform well (esp. "Quantum Monte Carlo" methods)

→ ex: Markov Chain Monte Carlo methods (Metropolis–Hastings algorithm, etc.)

A testbed: Stoquastic Hamiltonians

An Hamiltonian is stoquastic (a.k.a. "sign-free") if its off-diagonal terms are non-positive

Complexity of Stoquastic Frustration-Free Hamiltonians

Authors: **Sergey Bravyi** and **Barbara Terhal** | AUTHORS INFO & AFFILIATIONS

https://doi.org/10.1137/08072689X

Random walk that converges rapidly to the ground state distribution when Gapped + Stoquastic + Frustration-Free

Quantum
the open journal for quantum science

PAPERS PERSPECTIVE

A rapidly mixing Markov chain from any gapped quantum many-body system

Sergey Bravyi[1], Giuseppe Carleo[2], David Gosset[3,4], and Yinchen Liu[3,4]

Extended to any gapped Hamiltonian, when amplitude ratios $\langle x | \psi_H \rangle / \langle y | \psi_H \rangle$ are efficiently computable

A testbed: Stoquastic Hamiltonians

Our question

Does a stoquastic Hamiltonian exist whose ground state is easy to compute for quantum algorithms but classically hard, no matter what guiding state is given?

- Frustration-free Hamiltonians won't work
- We want to rule out all guiding states (including the exact one: $|\tilde{\psi}\rangle = |\psi_H\rangle$)
- Prior work has answered this question positively only for certain types of guiding states

(Sub)Exponential advantage of adiabatic Quantum computation with no sign problem

Authors: András Gilyén, Matthew B. Hastings, Umesh Vazirani | Authors Info & Claims

STOC 2021: Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing
Pages 1357 - 1369 • https://doi.org/10.1145/3406325.3451060

A testbed: Stoquastic Hamiltonians

Our question

> Does a stoquastic Hamiltonian exist whose ground state is easy to compute for quantum algorithms but classically hard, no matter what guiding state is given?

- We construct such a Hamiltonian as follows:

Bounded-degree graph G over 2^n vertices → Adjacency matrix A_G → Sparse stoquastic Hamiltonian $H = -A_G$ → poly(n) quantum alg. / $2^{\Omega(n^{1/4})}$ classical lower bound

Provided via an adjacency-list oracle

Ground state of H = Top eigenvector of A_G

Main result

Proof idea (on the board)

Conclusion

Sampling from the ground state distribution of a sparse stoquastic Hamiltonian

Requires (sub)exponentially more queries to H than quantumly

∃ A “natural” Hamiltonian problem that is much harder for classical computers (than for quantum computers), even when assisted with some “basic” quantum resources

Samples from the distribution of any guiding state

Open questions:

- A similar result for more physically-motivated Hamiltonians (e.g., locality constraints)?
- A larger classical/quantum separation (truly exponential speedup - $2^{\Omega(n)}$ classical lower bound)?
- Other graphs for which sampling the top eigenvector of the adjacency matrix is classically hard?