Dequantizing the Quantum Singular Value Transformation: Hardness and Applications to Quantum Chemistry and the Quantum PCP Conjecture

ArXiv: 2111.09079
STOC'22, QIP'22

Sevag Gharibian
Department of Computer Science
Paderborn University

François Le Gall
Graduate School of Mathematics
Nagoya University

SUSTech-Nagoya workshop on Quantum Science
June 3rd, 2022

1-Slide Overview of the Results

Main result:

computing an estimation of the ground state energy with inverse-polynomial precision, (when given a rough estimation of the ground state)

> We show that a central computational problem considered by quantum algorithms for quantum chemistry is BQP-complete.

"as hard as simulating universal polynomial-size quantum circuits"

➡ This gives theoretical foundations to claim the superiority of quantum algorithms for chemistry!

Second result:

> We show that computing an estimation of the ground state energy with constant precision can be done classically in polynomial time.

➡ This shows that the superiority of quantum algorithms comes from the improved precision achievable in the quantum setting

> To prove the second result, we show how to "dequantize" the Quantum Singular Value Transformation with constant precision

This dequantization result has implications to the famous quantum PCP conjecture, which is one of the central conjectures in quantum complexity theory

Quantum Chemistry and Eigenvalue Estimation

- ✓ Quantum chemistry is considered as one of the most promising applications of quantum computers
- ✓ From a computer science perspective as well, quantum chemistry is attractive since the main goal is clearly defined:

> compute a good estimation of the ground state energy of a local Hamiltonian representing the system
>
> (in more mathematical terms: compute a good estimation of the smallest eigenvalue of a sparse Hermitian matrix)

- ✓ The most rigorous approaches, first proposed by [Abrams, Lloyd 99] [Aspuru-Guzik, Dutoi, Love, Head-Gordon 05], are based on quantum phase estimation

  TODAY'S FOCUS

- ✓ Other promising approaches such as variational quantum algorithms are also actively studied but these approaches are mostly heuristic-based and their performance is thus much more difficult to evaluate in a rigorous way

Guided Local Hamiltonian Problem

Informal description:

> input: a sparse Hamiltonian H acting on n qubits
> a quantum state that has good overlap with the ground state of H
>
> output: an estimation of the ground state energy

Formal description:

GLH(s,ε,δ) "Guided local Hamiltonian problem"

$s \geq 1$: sparsity parameter
$\delta \in (0,1]$: overlap parameter
$\varepsilon \in (0,1]$: precision parameter

> input: an s-sparse Hamiltonian H acting on n qubits such that $\|H\| \leq 1$
> an n-qubit quantum state $|u\rangle$
>
> promise: $\|\Pi_H|u\rangle\| \geq \delta$
>
> output: an estimate $\tilde{\lambda}$ such that $|\tilde{\lambda} - \lambda_H| \leq \varepsilon$

H is s-sparse if it contains at most s non-zero entries per row and column (remember: H is a 2^n x 2^n matrix)

k-local ⇨ poly(n)2^k-sparse

λ_H: ground state energy of H (i.e., smallest eigenvalue)

Π_H: projection into the vector space spanned by the ground states of H

Guided Local Hamiltonian and Chemistry

GLH(s,ε,δ) "Guided local Hamiltonian problem"

input: an s-sparse Hamiltonian H acting on n qubits such that $\|H\| \leq 1$
an n-qubit quantum state $|u\rangle$

promise: $\|\Pi_H|u\rangle\| \geq \delta$

output: an estimate $\tilde{\lambda}$ such that $|\tilde{\lambda} - \lambda_H| \leq \varepsilon$

Guided Local Hamiltonians and Chemistry

Quantum-phase-estimation-based approach to quantum chemistry
(e.g., [Abrams, Lloyd 99] [Aspuru-Guzik, Dutoi, Love, Head-Gordon 05][Lee et al. 21])

1. Find a model for the chemical system (e.g., second quantization with finite-size basis), and express its Hamiltonian using qubits

   ➡ this gives a s-sparse Hamiltonian acting on n qubits, where s is polynomial in n (for instance $s = O(n^4)$ [Lee et al. 21], $s = O(n^2)$ [MacClean et al. 14])

2. Find a quantum state that has good overlap with the ground state

   ➡ the Hartree-Fock method typically recovers 99% of the total energy [Whitfield et al. 13]

3. Apply quantum phase estimation

   ➡ running time polynomial in s, $1/\delta$ and $1/\varepsilon$ (polynomial in n when s is polynomial in n, and δ, ε are inverse-polynomial in n)

worked out explicitly in the framework of the Quantum Singular Value Transformation [Gilyen et al.19] [Martyn et al. 21] and eigenstate filtering [Lin, Yu]

GLH(s,ε,δ) "Guided local Hamiltonian problem"

input: an s-sparse Hamiltonian H acting on n qubits such that $\|H\| \leq 1$
an n-qubit quantum state $|u\rangle$

promise: $\|\Pi_H|u\rangle\| \geq \delta$

output: an estimate $\tilde{\lambda}$ such that $|\tilde{\lambda} - \lambda_H| \leq \varepsilon$

Guided Local Hamiltonians and Chemistry

Theorem (expressing the phase-estimation approach to quantum chemistry)

For any $s \leq poly(n)$ and any $\delta,\varepsilon \geq 1/poly(n)$, the problem GLH(s,ε,δ) can be solved in poly(n)-time with a quantum computer.

basis), and express its Hamiltonian using qubits

➡ this gives a s-sparse Hamiltonian acting on n qubits, where s is polynomial in n (for instance $s = O(n^4)$ [Lee et al. 21], $s = O(n^2)$ [MacClean et al. 14])

2. Find a quantum state that has good overlap with the ground state

➡ the Hartree-Fock method typically recovers 99% of the total energy [Whitfield et al. 13]

3. Apply quantum phase estimation

worked out explicitly in the framework of the Quantum Singular Value Transformation [Gilyen et al.19] [Martyn et al. 21] and eigenstate filtering [Lin, Yu]

➡ running time polynomial in s, 1/δ and 1/ε (polynomial in n when s is polynomial in n, and δ, ε are inverse-polynomial in n)

GLH(s,ε,δ) “Guided local Hamiltonian problem”

input: an s-sparse Hamiltonian H acting on n qubits such that $\|H\| \leq 1$
an n-qubit quantum state $|u\rangle$

promise: $\|\Pi_H|u\rangle\| \geq \delta$

output: an estimate $\tilde{\lambda}$ such that $|\tilde{\lambda} - \lambda_H| \leq \varepsilon$

Guided Local Hamiltonian Problem

Theorem (expressing the phase-estimation approach to quantum chemistry)

> For any s ≤ poly(n) and any δ,ε ≥1/poly(n), the problem GLH(s,ε,δ) can be solved in poly(n)-time with a quantum computer.

Is it really a hard problem for classical computers?

Reasons why it may be hard (and counter-arguments):

- ✓ phase estimation solves integer factoring and many other hard problems
  - but is it really hard for a sparse matrix?
- ✓ if no guiding state $|u\rangle$ is given, then we know the problem is very hard (QMA-hard)
  - but having a guiding state significantly simplifies the problem…
  - (if $|u\rangle$ is exactly the ground state, then λ_H can be computed easily classically)

GLH(s,ε,δ) "Guided local Hamiltoni…

input: an s-sparse Hamiltonian H acti…
an n-qubit quantum state $|u\rangle$

promise: $\| \Pi_H |u\rangle \| \geq \delta$

output: an estimate $\tilde{\lambda}$ such that $| \tilde{\lambda} - \lambda_H |$…

Consider the 2^n-dimensional vector u (stored in a classical Random-Access-Memory)

Assume that $u_1 \neq 0$

easy to compute since H is sparse

Then $\lambda_H = \dfrac{\text{first coordinate of the vector } Hu}{\text{first coordinate of the vector } u}$

Our Results

n: number of qubits
(H: 2^n x 2^n matrix)

$s \geq 1$: number of non-zero entries in each row of H
$\delta \in (0,1]$: overlap between $|u\rangle$ and the ground state
$\varepsilon \in (0,1]$: precision parameter

Theorem (expressing the phase-estimation approach to quantum chemistry)

For any $s \leq poly(n)$ and any $\delta,\varepsilon \geq 1/poly(n)$, the problem GLH(s,ε,δ) can be solved in poly(n)-time with a quantum computer.

Our first result

The problem GLH(s,ε,δ) is BQP-hard for $s = poly(n)$, $\varepsilon = 1/poly(n)$ and $\delta = 1/2$.

"If there exists a classical algorithm that solves GLH(s,ε,δ) with inverse-polynomial precision (even for δ = 1/2), then any quantum polynomial time computation (e.g., Shor algorithm) can be simulated classically in polynomial time."

The holy grail in quantum chemistry is to get estimation of the ground state energy with precision less than the "chemical accuracy" (about 1.6 millihartree), which corresponds to inverse-polynomial precision after normalizing the Hamiltonian

➡ This gives some theoretical foundations to claim the superiority of quantum algorithms for chemistry

Our second result

For any $s \leq poly(n)$ and any constant $\delta,\varepsilon > 0$, the problem GLH(s,ε,δ) can be solved in poly(n)-time with a classical computer.

"The problem can be solved classically in polynomial time with constant precision even with arbitrarily small constant overlap δ."

➡ This shows that the superiority of quantum algorithms comes from the improved precision achievable in the quantum setting

Proof of Hardness

to simplify the explanations I will assume that the quantum circuit does not make any error (i.e., on each input x, the circuit correctly outputs "yes" with probability 1 or outputs "no" with probability 1)

Our first result

The problem GLH(s,ε,δ) is BQP-hard for s = poly(n), ε = 1/poly(n) and δ = 1/2.

We show that if we can solve efficiently this problem, we can efficiently solve any decision problem that can be solved by a polynomial-size quantum circuit.

✓ Consider a polynomial-size quantum circuit $U = U_m \ldots U_1$ with m = poly(n)

From this circuit, we need to show how to efficiently classically create a sparse Hamiltonian H and a state $|u\rangle$ with $\|\Pi_H|u\rangle\| \geq \delta$ such that a solution to GLH(s,ε,δ) gives information about the output of the circuit.

GLH(s,ε,δ) "Guided local Hamiltonian problem"

input: an s-sparse Hamiltonian H acting on n qubits such that $\|H\| \leq 1$
an n-qubit quantum state $|u\rangle$

promise: $\|\Pi_H|u\rangle\| \geq \delta$

output: an estimate $\tilde{\lambda}$ such that $|\tilde{\lambda} - \lambda_H| \leq \varepsilon$

Proof of Hardness

to simplify the explanations I will assume that the quantum circuit does not make any error (i.e., on each input x, the circuit correctly outputs "yes" with probability 1 or outputs "no" with probability 1)

Our first result

The problem GLH(s,ε,δ) is BQP-hard for s = poly(n), ε = 1/poly(n) and δ = 1/2.

We show that if we can solve efficiently this problem, we can efficiently solve any decision problem that can be solved by a polynomial-size quantum circuit.

- ✓ Consider a polynomial-size quantum circuit $U = U_m \ldots U_1$ with m = poly(n)
- ✓ We apply Kitaev's circuit-to-Hamiltonian construction [Kitaev et al. 02, 06] to map the circuit U to a 5-local Hamiltonian

$$H = H_{in} + H_{prop} + H_{out} + H_{stab}$$

so that the ground space of H "simulates" U. In particular, we can show that if U outputs "yes" on the input then $\lambda_H = 0$, while if U outputs "no" on the input then $\lambda_H \geq 1/m^3 = 1/\mathrm{poly}(n)$.

inverse-polynomial gap

- ✓ In the first case ($\lambda_H = 0$) we know that the ground state is the "history state"

$$|\psi_{\mathrm{hist}}\rangle = \frac{1}{\sqrt{m+1}} \sum_{t=0}^{m} U_t \cdots U_1 |x\rangle_A |0\cdots 0\rangle_B |t\rangle_C \, .$$

x: input of U

input space ($|x\rangle_A$) work space ($|0\cdots 0\rangle_B$) clock space ($|t\rangle_C$)

the state $|u\rangle = |x\rangle_A\, |0\cdots 0\rangle_B\, |0\rangle_C$, which is easy to generate classically, has non-trivial overlap with $|\psi_{\mathrm{hist}}\rangle$

Problem: in the second case ($\lambda_H \geq 1/m^3$), how to generate efficiently a state $|u\rangle$ that has good overlap with the ground space of H? (we don't even have a good mathematical description of this space!)

Consider the 6-local Hamiltonian $H' = H \otimes |0\rangle_D\langle 0|_D + \frac{I}{2m^3} \otimes |1\rangle_D\langle 1|_D$

- If U outputs "yes" then $\lambda_{H'} = 0$. The corresponding ground state is $|\psi_{\text{hist}}\rangle\, |0\rangle_D$.
- If U outputs "no" then $\lambda_{H'} = 1/(2m^3) = 1/\text{poly}(n)$.
  Any state of the form $|\varphi\rangle\, |1\rangle_D$ is a ground state (for any $|\varphi\rangle$).

the state $|u'\rangle = |x\rangle_A\, |0 \cdots 0\rangle_B\, |0\rangle_C\, |+\rangle_D$ has non-trivial overlap even for the case where U outputs "no" !

$$H = H_{in} + H_{prop} + H_{out} + H_{stab}$$

so that the ground space of H "simulates" U. In particular, we can show that if U outputs "yes" on the input then $\lambda_H = 0$, while if U outputs "no" on the input then $\lambda_H \geq 1/m^3 = 1/\text{poly}(n)$.

✓ In

We have constructed a 6-local Hamiltonian (and thus s-sparse with s = poly(n)) and a state $|u\rangle$ with overlap $\delta \geq 1/\text{poly}(n)$ such that solving GLH(s,ε,δ) with ε = 1/poly(n) identifies the output of the circuit U on the input x.

input space work space clock space

the state $|u\rangle = |x\rangle_A\, |0 \cdots 0\rangle_B\, |0\rangle_C$, which is easy to generate classically, has non-trivial overlap with $|\psi_{\text{hist}}\rangle$

Problem: in the second case ($\lambda_H \geq 1/m^3$), how to generate efficiently a state $|u\rangle$ that has good overlap with the ground space of H? (we don't even have a good mathematical description of this space!)

Proof of Hardness

to simplify the explanations I will assume that the quantum circuit does not make any error (i.e., on each input x, the circuit correctly outputs “yes” with probability 1 or outputs “no” with probability 1)

Our first result

The problem GLH(s,ε,δ) is BQP-hard for s = poly(n), ε = 1/poly(n) and δ = 1/2.

We show that if we can solve efficiently this problem, we can efficiently solve any decision problem that can be solved by a polynomial-size quantum circuit.

✓ Consider a polynomial-size quantum circuit $U = U_m \dots U_1$ with m = poly(n)

We have constructed a 6-local Hamiltonian (and thus s-sparse with s = poly(n)) and a state $|u\rangle$ with overlap δ ≥ 1/poly(n) such that solving GLH(s,ε,δ) with ε = 1/poly(n) identifies the output of the circuit U on the input x.

Remains to do:

✓ increase the overlap δ to 1/2 (we use “pre-idling”)

✓ deal with the case where the quantum circuit makes errors

Proof of Hardness

to simplify the explanations I will assume that the quantum circuit does not make any error (i.e., on each input x, the circuit correctly outputs "yes" with probability 1 or outputs "no" with probability 1)

Our first result

The problem GLH(s,ε,δ) is BQP-hard for s = poly(n), ε = 1/poly(n) and δ = 1/2.

We show that if we can solve efficiently this problem, we can efficiently solve any decision problem that can be solved by a polynomial-size quantum circuit.

✓ Consider a polynomial-size quantum circuit $U = U_m \dots U_1$ with m = poly(n)

We have constructed a 6-local Hamiltonian (and thus s-sparse with s = poly(n)) and a state $|u\rangle$ with overlap δ ≥ 1/poly(n) such that solving GLH(s,ε,δ) with ε = 1/poly(n) identifies the output of the circuit U on the input x.

Remains to do:

✓ increase the overlap δ to 1/2 (we use "pre-idling")

✓ deal wit[h] makes e[...]

We replace $U = U_m \dots U_1$ by $U = U_m \dots U_1 \underbrace{I \dots I}_{\text{poly(n) times}}$

Then a large part of the history state becomes trivial

Proof of Hardness: Open Problems

Our first result

The problem GLH(s,ε,δ) is BQP-hard for $s = \text{poly}(n)$, $\varepsilon = 1/\text{poly}(n)$ and $\delta = 1/2$.

The 6-local Hamiltonian H used to prove the hardness encodes the computation of an arbitrary quantum circuit

open problem #1: improve the parameters

$\delta \rightarrow 1\text{-}1/\text{poly}(n)$, better sparsity (e.g., 2-local Hamiltonian)

open problem #2: prove the hardness for the Hamiltonians occurring in quantum chemistry

GLH(s,ε,δ) "Guided local Hamiltonian problem"

input: an s-sparse Hamiltonian H acting on n qubits such that $\|H\| \leq 1$
an n-qubit quantum state $|u\rangle$

promise: $\|\Pi_H|u\rangle\| \geq \delta$

output: an estimate $\tilde{\lambda}$ such that $|\tilde{\lambda} - \lambda_H| \leq \varepsilon$

Second Result

Our first result

The problem GLH(s,ε,δ) is BQP-hard for $s = \text{poly}(n)$, $\varepsilon = 1/\text{poly}(n)$ and $\delta = 1/2$.

Our second result

For any $s \leq \text{poly}(n)$ and any constant $\delta,\varepsilon > 0$, the problem GLH(s,ε,δ) can be solved in poly(n)-time with a classical computer.

"The problem can be solved classically in polynomial time with constant precision even with arbitrarily small constant overlap δ."

concretely, we assume that we can perform ℓ_2-sampling from u as in prior works in dequantization [Tang 19][Chia et al. 20](see also [Van den Nest 10]):

one sample gives (i,u_i) with probability $|u_i|^2$

GLH(s,ε,δ) "Guided local Hamiltonian problem"

input: an s-sparse Hamiltonian H acting on n qubits such that $\|H\| \leq 1$

~~an n-qubit quantum state $|u\rangle$~~ an efficient classical representation of a unit-norm vector $|u\rangle \in \mathbb{C}^{2^n}$

promise: $\|\Pi_H|u\rangle\| \geq \delta$

output: an estimate $\tilde{\lambda}$ such that $|\tilde{\lambda} - \lambda_H| \leq \varepsilon$

Second Result: Dequantizing the QSVT

"Decision version" of the guided local Hamiltonian problem (for a < b)

input: an s-sparse Hamiltonian H acting on n qubits such that $\|H\| \leq 1$
ℓ_2-sampling access to a unit-norm vector $|u\rangle \in \mathbb{C}^{2^n}$
promise: $\|\Pi_H|u\rangle\| \geq \delta$
either $\lambda_H \leq a$ or $\lambda_H \geq b$ holds
goal: decide which of $\lambda_H \leq a$ or $\lambda_H \geq b$ holds

Eigenvalue decomposition: $H = \sum_{i=1}^{2^n} \sigma_i |v_i\rangle\langle v_i|$ with $-1 \leq \sigma_1 \leq \sigma_2 \leq \ldots \leq \sigma_{2^n} \leq 1$
(we have $\lambda_H = \sigma_1$)

For any polynomial p: $p(H) = \sum_{i=1}^{2^n} p(\sigma_i) |v_i\rangle\langle v_i|$

QSVT: given a qubitization of H, compute a qubitization of p(H)

- ✓ Definition of the framework and quantum algorithms: [Gilyén, Su, Low and Wiebe 19] [Low and Chuang 17,19] [Martyn, Rossi, Tan and Chuang 21],…
- ✓ Dequantization possible for low-rank matrices: [Chia, Gilyén, Li, Lin, Tang and Wang 20])

Second Result: Dequantizing the QSVT

"Decision version" of the guided local Hamiltonian problem (for $a < b$)

input: an s-sparse Hamiltonian H acting on n qubits such that $\|H\| \leq 1$
ℓ_2-sampling access to a unit-norm vector $|u\rangle \in \mathbb{C}^{2^n}$

promise: $\|\Pi_H|u\rangle\| \geq \delta$
either $\lambda_H \leq a$ or $\lambda_H \geq b$ holds

goal: decide which of $\lambda_H \leq a$ or $\lambda_H \geq b$ holds

Our second result

For any $s \leq poly(n)$ and any constant $\delta > 0$ and any constants $a<b$ the problem can be solved in poly(n)-time with a classical computer.

$. \leq \sigma_{2^n} \leq 1$
$\sigma_1)$

Lemma:
[Low, Chuang 19]

There exists a polynomial q of degree $O(1/(b-a))$ such that $q(x) \in [0,1]$ for all $x \in [-1,1]$, $q(x) \approx 1$ if $x \leq a$ and $q(x) \approx 0$ if $x \geq b$

q(H) is $O(s^{O(1/(b-a))})$ -sparse,
ℓ_2-sampling access to u

$\Rightarrow$ $\begin{cases} \|q(H)|u\rangle\| \approx 0 \text{ if } \lambda_H \geq b \\ \|q(H)|u\rangle\| \geq \delta q(\lambda_H) \approx \delta \text{ if } \lambda_H \leq a \end{cases}$

QSVT distinguishes the two cases in poly(s,1/δ,1/(b-a)) time

We show that classically, this can be done in $O(s^{O(1/(b-a))})$ time (for δ constant)

Quantum PCP Conjecture

no guiding vector!

LH(k,a,b) "Local Hamiltonian problem" (for $a < b$)

input: a k-local Hamiltonian H acting on n qubits such that $\|H\| \leq 1$
promise: either $\lambda_H \leq a$ or $\lambda_H \geq b$ holds
goal: decide which of $\lambda_H \leq a$ or $\lambda_H \geq b$ holds

H is k-local if it can be written as a sum of poly(n) terms, where each term acts on at most k qubits

k-local ⇨ poly(n)2^k-sparse

known: [Kitaev et al. 02,06]

There exist $a,b \in [-1,1]$ with $b-a = 1/\text{poly}(n)$ such that LH(2,a,b) is QMA-hard.

Quantum generalization of the class NP

"there exist local Hamiltonians for which estimating the ground energy with inverse-polynomial precision is very hard"

Quantum PCP conjecture:

There exist $k=O(1)$ and $a,b \in [-1,1]$ with $b-a = \Omega(1)$ such that LH(k,a,b) is QMA-hard.

"there exist local Hamiltonians for which estimating the ground energy even with constant precision is very hard"

Our Result

no guiding vector!

LHS(k,a,b) "Local Hamiltonian problem with samplable state" (for $a < b$)

input: a k-local Hamiltonian H acting on n qubits such that $\|H\| \leq 1$

promise: either $\lambda_H \leq a$ or $\lambda_H \geq b$ holds

there exists an efficiently-samplable state $|u\rangle$ such that $\|\Pi_H|u\rangle\| = \Omega(1)$

output: goal: decide which of $\lambda_H \leq a$ or $\lambda_H \geq b$ holds

there exists a classical description of $|u\rangle$ such that (approximate) ℓ_2-sampling can be done in poly time

Our result: For any $k = O(\log n)$ and any $a,b \in [-1,1]$ with $b-a = \Omega(1)$, LHS(k,a,b) is not QMA-hard (unless QMA=MA).

"unless the Quantum PCP conjecture is false, Hamiltonians involved in the Quantum PCP conjecture do not have a non-trivial approximation of their ground state by efficiently-samplable state"

Proof: We show that LHS(k,a,b) is in MA

The classical prover simply guesses the classical description of the state $|u\rangle$, and the classical verifier applies our dequantized version of the QSVT to check which of $\lambda_H \leq a$ or $\lambda_H \geq b$ holds

Conclusion

Main result:

computing an estimation of the ground state energy with inverse-polynomial precision (given a rough estimation of the ground state)

> We show that a central computational problem considered by quantum algorithms for quantum chemistry is BQP-complete.

"as hard as simulating universal polynomial-size quantum circuits"

➡ This gives theoretical foundations to claim the superiority of quantum algorithms for chemistry!

Second result:

> We show that computing an estimation of the ground state energy with constant precision can be done classically in polynomial time.

➡ This shows that the superiority of quantum algorithms comes from the improved precision achievable in the quantum setting

> To prove the second result, we show how to "dequantize" the Quantum Singular Value Transformation (for a constant-degree polynomial) with constant precision

This dequantization result gives a new perspective on the famous quantum PCP conjecture, which is one of the central conjectures in quantum complexity theory

Open Problems

Main result:

computing an estimation of the ground state energy with inverse-polynomial precision (given a rough estimation of the ground state)

We show that a central computational problem considered by quantum algorithms for quantum chemistry is BQP-complete.

"as hard as simulating universal polynomial-size quantum circuits"

➡ This gives theoretical foundations to claim the superiority of quantum algorithms for chemistry!

open problem #1: improve the parameters

$\delta \rightarrow$ 1-1/poly(n), better sparsity (e.g., 2-local Hamiltonian)

open problem #2: prove the hardness for the Hamiltonians occurring in quantum chemistry

open problem #3: give theoretical foundations for the approaches based on variational quantum algorithms for quantum chemistry