The optimal depth of variational quantum algorithms is QCMA-hard to approximate[1]

Lennart Bittel Sevag Gharibian Martin Kliesch

[1]Institute for Theoretical Physics
Heinrich Heine University Düsseldorf
Germany

[2]Department of Computer Science
Inst. for Photonic Quantum Systems
Paderborn University
Germany

[3]Heinrich Heine University Düsseldorf
Hamburg University of Technology
Germany

PADERBORN UNIVERSITY

PhoQS

[1]arXiv:2211.12519

Outline

Germany to invest €2B in quantum technologies

11 May 2021 | News

In one of the biggest spending plans of its kind in the world, the government commits to develop country's first quantum computer

By Éanna Kelly

German Science Minister, Anja Karliczek Photo: Anja Karliczek website

Germany is to invest €2 billion in quantum computing and related technologies over five years, under a plan that dwarfs that of almost every other country, with the education and research ministry committing €1.1 billion by 2025 for R&D, while the economy ministry will contribute €878 million to develop applications.

The German Aerospace Centre will get most of the money, some €740 million, to team up with industry.

Announcing the plan on Tuesday, science minister Anja Karliczek, said the government aims to build a competitive quantum computer in five years, while growing a network of companies to develop applications.

Noisy Intermediate-Scale Quantum (NISQ) computation era

Question: What can we do with near-term quantum devices with

- small number of qubits,
- short circuit depth,
- limited connectivity between qubits?

Variational Quantum Algorithms (VQA)

Idea: Hybrid classical-quantum setup

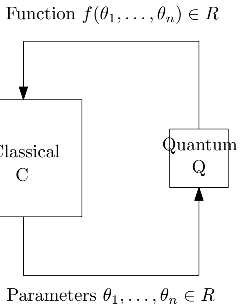

Goals

- Wish to minimize some cost function $f : \mathbb{R}^n \to \mathbb{R}$
- Variationally choose parameters $\theta_i \in \mathbb{R}$ (via gradient descent, machine learning, etc)
- Our focus: Keep Q as small as possible (few qubits, low depth, etc)

Under the hood

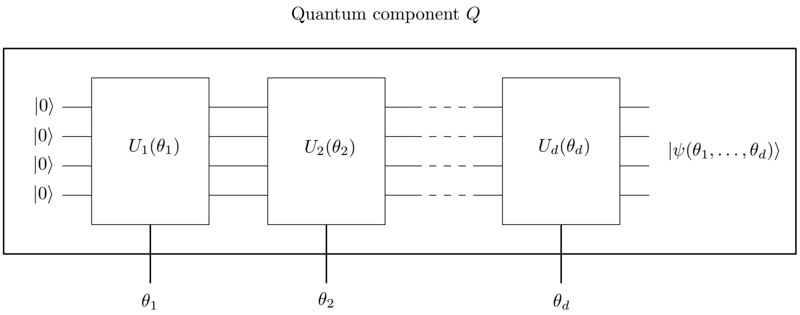

- Given set of Hamiltonians $\{H_k\}$, choose unitaries $U_k(\theta_k) = e^{i\theta_k H_k}$ for $k = 1, \ldots, d$
- Roughly, a "fast-forwarded" version of standard Trotterization of Hamiltonian evolution

Under the hood

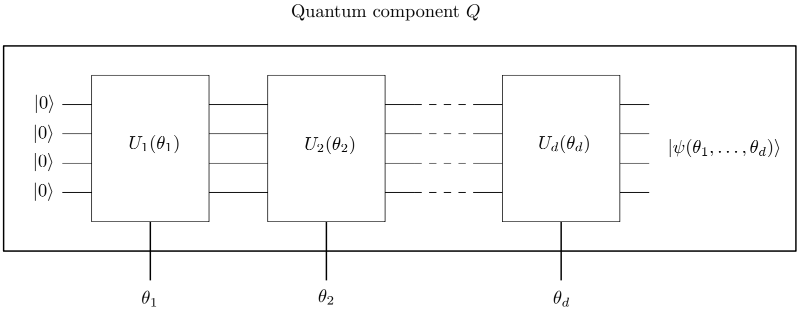

- Given set of Hamiltonians $\{H_k\}$, choose unitaries $U_k(\theta_k) = e^{i\theta_k H_k}$ for $k = 1,\ldots,d$
- Roughly, a "fast-forwarded" version of standard Trotterization of Hamiltonian evolution

Goal:

- Minimize "depth" d, i.e. number of rotations applied
- Crucial for NISQ devices: Low depth $\Rightarrow$ circuit completes before noise destroys computation

Quantum Approximate Optimization Algorithm (QAOA)

- Introduced in 2014 by Farhi, Goldstone, Gutmann
- Tries to approximately solve hard combinatorial problems, e.g. MAX CUT

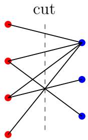

For MAX CUT, alternate application of $H_1 = \sum_{\text{edges } (i,j)} Z_i \otimes Z_j$ and $H_2 = \sum_i X_i$.

Quantum Approximate Optimization Algorithm (QAOA)

- Introduced in 2014 by Farhi, Goldstone, Gutmann
- Tries to approximately solve hard combinatorial problems, e.g. MAX CUT

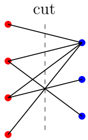

For MAX CUT, alternate application of $H_1 = \sum_{\text{edges }(i,j)} Z_i \otimes Z_j$ and $H_2 = \sum_i X_i$.

Algorithm

1. "Pick" variational angles $\theta_1, \ldots, \theta_d$.
2. Use Q to prepare state $|\psi\rangle = e^{i\theta_d H_2} e^{i\theta_d H_1} \cdots e^{i\theta_2 H_2} e^{i\theta_1 H_1} |+\cdots+\rangle$.
3. Measure $|\psi\rangle$ in standard basis to get string x, which defines a cut in graph.

Quantum Approximate Optimization Algorithm (QAOA)

- Introduced in 2014 by Farhi, Goldstone, Gutmann
- Tries to approximately solve hard combinatorial problems, e.g. MAX CUT

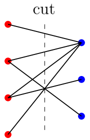

For MAX CUT, alternate application of $H_1 = \sum_{\text{edges }(i,j)} Z_i \otimes Z_j$ and $H_2 = \sum_i X_i$.

Algorithm

1. "Pick" variational angles $\theta_1, \ldots, \theta_d$.
2. Use Q to prepare state $|\psi\rangle = e^{i\theta_d H_2} e^{i\theta_d H_1} \cdots e^{i\theta_2 H_2} e^{i\theta_1 H_1} |+\cdots+\rangle$.
3. Measure $|\psi\rangle$ in standard basis to get string x, which defines a cut in graph.

Question: What is the "right" depth d to use?

(Selected) previous work

Early good news:

- Level-1 QAOA achieves 0.6924-approximation for MAX CUT [Farhi, Goldstone, Gutmann 2014]

(Selected) previous work

Early good news:

- Level-1 QAOA achieves 0.6924-approximation for MAX CUT [Farhi, Goldstone, Gutmann 2014]
- Level-1 QAOA's output distribution cannot be efficiently simulated classically [Farhi, Harrow 2016]

(Selected) previous work

Early good news:

- Level-1 QAOA achieves 0.6924-approximation for MAX CUT [Farhi, Goldstone, Gutmann 2014]
- Level-1 QAOA's output distribution cannot be efficiently simulated classically [Farhi, Harrow 2016]
- QAOA with poly depth is universal [Lloyd 2018]

(Selected) previous work

Early good news:

- Level-1 QAOA achieves 0.6924-approximation for MAX CUT [Farhi, Goldstone, Gutmann 2014]
- Level-1 QAOA's output distribution cannot be efficiently simulated classically [Farhi, Harrow 2016]
- QAOA with poly depth is universal [Lloyd 2018]

Later, less good news:

- $O(1)$-level QAOA cannot outperform Goemans-Williamson algorithm for MAX CUT [Bravyi, A. Kliesch, Koenig, Tang 2020]

(Selected) previous work

Early good news:

- Level-1 QAOA achieves 0.6924-approximation for MAX CUT [Farhi, Goldstone, Gutmann 2014]
- Level-1 QAOA's output distribution cannot be efficiently simulated classically [Farhi, Harrow 2016]
- QAOA with poly depth is universal [Lloyd 2018]

Later, less good news:

- $O(1)$-level QAOA cannot outperform Goemans-Williamson algorithm for MAX CUT [Bravyi, A. Kliesch, Koenig, Tang 2020]
- Many (many) heuristic studies suggest hard to optimize parameters (e.g. barren plateaus)

(Selected) previous work

Early good news:

- Level-1 QAOA achieves 0.6924-approximation for MAX CUT [Farhi, Goldstone, Gutmann 2014]
- Level-1 QAOA's output distribution cannot be efficiently simulated classically [Farhi, Harrow 2016]
- QAOA with poly depth is universal [Lloyd 2018]

Later, less good news:

- $O(1)$-level QAOA cannot outperform Goemans-Williamson algorithm for MAX CUT [Bravyi, A. Kliesch, Koenig, Tang 2020]
- Many (many) heuristic studies suggest hard to optimize parameters (e.g. barren plateaus)
- NP-hard to optimize angles θ_k if Hamiltonian sequence $(H_1, \ldots, H_d)$ and depth d prespecified [Bittel, M. Kliesch, 2021]

(Selected) previous work

Early good news:

- Level-1 QAOA achieves 0.6924-approximation for MAX CUT [Farhi, Goldstone, Gutmann 2014]
- Level-1 QAOA's output distribution cannot be efficiently simulated classically [Farhi, Harrow 2016]
- QAOA with poly depth is universal [Lloyd 2018]

Later, less good news:

- $O(1)$-level QAOA cannot outperform Goemans-Williamson algorithm for MAX CUT [Bravyi, A. Kliesch, Koenig, Tang 2020]
- Many (many) heuristic studies suggest hard to optimize parameters (e.g. barren plateaus)
- NP-hard to optimize angles θ_k if Hamiltonian sequence $(H_1, \ldots, H_d)$ and depth d prespecified [Bittel, M. Kliesch, 2021]

This work: How hard to estimate the optimal depth, d, for VQA/QAOA?

Definition of VQA minimization used by [Bittel, M. Kliesch, 2021]

Recall:

- NP-hard to optimize angles θ_k if Hamiltonian sequence $(H_1, \ldots, H_d)$ and depth d prespecified [Bittel, M. Kliesch, 2021]

VQA minimization (MIN-VQA) [Bittel, M. Kliesch, 2021]

- Input: Sequence $(H_1, \ldots, H_L)$ of local Hamiltonians, observable M
- Output: Angles $(\theta_1, \ldots, \theta_L)$ such that $|\psi\rangle := e^{i\theta_L G_L} \cdots e^{i\theta_1 G_1} |0 \cdots 0\rangle$ minimizes $\langle\psi|M|\psi\rangle$.

Definition of VQA minimization used by [Bittel, M. Kliesch, 2021]

Recall:

- NP-hard to optimize angles θ_k if Hamiltonian sequence $(H_1, \ldots, H_d)$ and depth d prespecified [Bittel, M. Kliesch, 2021]

VQA minimization (MIN-VQA) [Bittel, M. Kliesch, 2021]

- Input: Sequence $(H_1, \ldots, H_L)$ of local Hamiltonians, observable M
- Output: Angles $(\theta_1, \ldots, \theta_L)$ such that $|\psi\rangle := e^{i\theta_L G_L} \cdots e^{i\theta_1 G_1} |0 \cdots 0\rangle$ minimizes $\langle\psi|M|\psi\rangle$.

In words:

- Rotation axes (i.e. Hamiltonians) and their order of application fixed
- Implicitly, this also fixes the depth L of the ansatz
- Question: What if we relax these restrictions, and focus purely on depth minimization?

Outline

Formalizing depth minimization

VQA minimization (MIN-VQA)

- Input: Set H of local Hamiltonians, observable M, depth thresholds $d_1 \leq d_2$
- Output:

YES: if $\exists$ at most d_1 angles $(\theta_1, \ldots, \theta_{d_1}) \in \mathbb{R}^{d_1}$ and Hamiltonians $(G_1, \ldots, G_{d_1}) \in H^{\times d_1}$ s.t.

$$|\psi\rangle := e^{i\theta_{d_1} G_{d_1}} \cdots e^{i\theta_1 G_1} |0 \cdots 0\rangle \qquad \text{satisfies } \langle\psi|M|\psi\rangle \leq 1/3.$$

NO: if $\forall$ sequences of at most d_2 angles $(\theta_1, \ldots, \theta_{d_2}) \in \mathbb{R}^{d_2}$ and $(G_1, \ldots, G_{d_2}) \in H^{\times d_2}$,

$$|\psi\rangle := e^{i\theta_{d_2} G_{d_2}} \cdots e^{i\theta_1 G_1} |0 \cdots 0\rangle \qquad \text{satisfies } \langle\psi|M|\psi\rangle \geq 2/3.$$

Notes:

- Closer to definition of ADAPT-VQE [Grimsley, Economou, Barnes, Mayhall 2019]
- Containment in QCMA straightforward: prover sends (θ_i) and (G_i), verifier runs Hamiltonian simulation.

Our result for MIN-VQA

Theorem 1

For any $\epsilon > 0$, it is QCMA-hard to distinguish between the YES and NO cases of MIN-VQA, even if

$$\frac{d_2}{d_1} \geq N^{1-\epsilon},$$

for N the encoding size of the instance.

In words:

- Approximating optimal depth of VQA, even up to large multiplicative factors, is intractable
- First natural QCMA-hard to approximate problem
- (Aside: NP $\subseteq$ MA $\subseteq$ QCMA $\subseteq$ QMA.)

As for QAOA

QAOA minimization (MIN-QAOA)

- Input:
  - Set $H = \{H_b, H_c\}$ of local Hamiltonians
  - Quantum circuit preparing ground state $|\mathrm{gs}_b\rangle$ of H_b
  - Depth thresholds $d_1 \leq d_2$
- Output:

YES: if $\exists$ at most d_1 angles $(\theta_1, \ldots, \theta_{d_1}) \in \mathbb{R}^{d_1}$ s.t.

$$|\psi\rangle := e^{i\theta(d_1)H_b} e^{i\theta(d_1-1)H_c} \cdots e^{i\theta_2 H_b} e^{i\theta_1 H_c} |\mathrm{gs}_b\rangle \qquad \text{satisfies } \langle\psi|M|\psi\rangle \leq 1/3.$$

NO: if $\forall$ sequences of at most d_2 angles $(\theta_1, \ldots, \theta_{d_2}) \in \mathbb{R}^{d_2}$,

$$|\psi\rangle := e^{i\theta(d_1)H_b} e^{i\theta(d_1-1)H_c} \cdots e^{i\theta_2 H_b} e^{i\theta_1 H_c} |\mathrm{gs}_b\rangle \qquad \text{satisfies } \langle\psi|M|\psi\rangle \geq 2/3.$$

Our result for MIN-QAOA

Theorem 2

For any $\epsilon > 0$, it is QCMA-hard to distinguish between the YES and NO cases of MIN-QAOA, even if

$$\frac{d_2}{d_1} \geq N^{1-\epsilon},$$

for N the encoding size of the instance.

Our result for MIN-QAOA

Theorem 2

For any $\epsilon > 0$, it is QCMA-hard to distinguish between the YES and NO cases of MIN-QAOA, even if

$$\frac{d_2}{d_1} \geq N^{1-\epsilon},$$

for N the encoding size of the instance.

Disclaimers:

- Assume “perfect”, idealized quantum computer (i.e. no noise, perfect gates, platform-independent, etc)
- Complexity results are worst-case, i.e. in practice special instances of problems might be easier to solve

Outline

The *de facto* "quantum NP"

Quantum Merlin-Arthur (QMA)

Promise problem $\mathbb{A} = (A_{\text{yes}}, A_{\text{no}}) \in$ QMA if $\exists$ poly-time uniformly generated quantum circuit family $\{Q_n\}$ s.t.:

- (YES case) If $x \in A_{\text{yes}}$, $\exists$ proof $|\psi_{\text{proof}}\rangle \in (\mathbb{C}^2)^{\otimes \text{poly}(n)}$, such that Q_n accepts with probability at least $2/3$.
- (NO case) If $x \in A_{\text{no}}$, then $\forall$ proofs $|\psi_{\text{proof}}\rangle \in (\mathbb{C}^2)^{\otimes \text{poly}(n)}$, Q_n accepts with probability at most $1/3$.

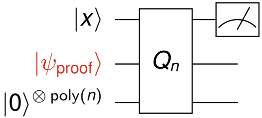

Wait... there's more than one definition "quantum NP"?

Wait... there's more than one definition "quantum NP"?

Wait... there's more than one definition of "quantum NP"?

Named after Snow White's dwarves:

1. (Doc) QMA
2. (Bashful) QMA_1: QMA with perfect completeness
3. (Happy) QCMA: QMA with classical proof
4. (Grumpy) $\mathrm{QMA}(2)$: QMA with "unentangled" proof of form $|\psi_1\rangle \otimes |\psi_2\rangle$
5. (Sneezy) NQP: Quantum TM accepts $x \in A_{\mathrm{yes}}$ in poly-time with probability > 0.
   (Equals $\mathrm{coC_{=}P}$ [Fenner, Green, Homer, Pruim, 1998].)
6. (Dopey) StoqMA: QMA with $\{|0\rangle, |+\rangle\}$ ancillae, classical gates, measurement in X basis

Quantum-Classical Merlin-Arthur (QCMA)

Promise problem $\mathbb{A} = (A_{\text{yes}}, A_{\text{no}}) \in \text{QCMA}$ if $\exists$ poly-time uniformly generated quantum circuit family $\{Q_n\}$ s.t.:

- (YES case) If $x \in A_{\text{yes}}$, $\exists$ proof $y \in \{0,1\}^{\text{poly}(n)}$, such that Q_n accepts with probability at least $2/3$.
- (NO case) If $x \in A_{\text{no}}$, then $\forall$ proofs $y \in \{0,1\}^{\text{poly}(n)}$, Q_n accepts with probability at most $1/3$.

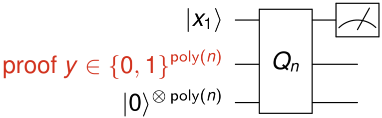

Question: What good is a classical proof to a quantum verifier?

Recall

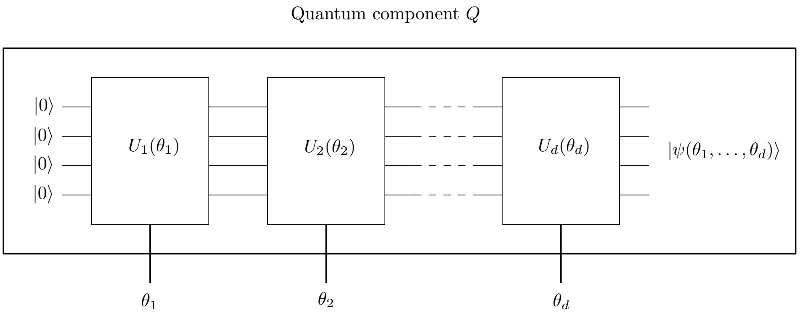

- Given set of Hamiltonians $\{H_k\}$, choose unitaries $U_k(\theta_k) = e^{i\theta_k H_k}$ for $k = 1, \ldots, d$
- Roughly, a "fast-forwarded" version of standard Trotterization of Hamiltonian evolution

Goal:

- Minimize "depth" d, i.e. number of rotations applied
- Crucial for NISQ devices: Low depth $\Rightarrow$ circuit completes before noise destroys computation

Outline

1. Variational Quantum Algorithms (VQAs)
2. Our results
3. Quantum Classical Merlin-Arthur (QCMA)
4. Proof sketches

Goal and challenges

Goal: Map given QCMA circuit V to instance (H, d, d') of MIN-VQA s.t. $\frac{d'}{d} \geq N^{1-\epsilon}$, and

$$\exists \text{ proof } y \text{ accepted by } V \implies \leq d \text{ VQA levels suffice to get "good" measurement result}$$
$$\forall \text{ proofs } y,\ V \text{ rejects} \implies > d' \text{ VQA levels required to get "good" measurement result}$$

Goal and challenges

Goal: Map given QCMA circuit V to instance (H, d, d') of MIN-VQA s.t. $\frac{d'}{d} \geq N^{1-\epsilon}$, and

$$\exists \text{ proof } y \text{ accepted by } V \implies \leq d \text{ VQA levels suffice to get "good" measurement result}$$
$$\forall \text{ proofs } y,\ V \text{ rejects} \implies > d' \text{ VQA levels required to get "good" measurement result}$$

Challenges:

1. Where will hardness of approximation (i.e. large ratio d'/d) come from?

Goal and challenges

Goal: Map given QCMA circuit V to instance (H, d, d') of MIN-VQA s.t. $\frac{d'}{d} \geq N^{1-\epsilon}$, and

$$\exists \text{ proof } y \text{ accepted by } V \implies \leq d \text{ VQA levels suffice to get "good" measurement result}$$
$$\forall \text{ proofs } y,\ V \text{ rejects} \implies > d' \text{ VQA levels required to get "good" measurement result}$$

Challenges:

1. Where will hardness of approximation (i.e. large ratio d'/d) come from?
2. MIN-VQA does not restrict which Hamiltonians are applied, in which order, with which rotation angles. How to enforce computational structure?

Goal and challenges

Goal: Map given QCMA circuit V to instance (H, d, d') of MIN-VQA s.t. $\frac{d'}{d} \geq N^{1-\epsilon}$, and

$$\exists \text{ proof } y \text{ accepted by } V \implies \leq d \text{ VQA levels suffice to get "good" measurement result}$$

$$\forall \text{ proofs } y, V \text{ rejects} \implies > d' \text{ VQA levels required to get "good" measurement result}$$

Challenges:

1. Where will hardness of approximation (i.e. large ratio d'/d) come from?
2. MIN-VQA does not restrict which Hamiltonians are applied, in which order, with which rotation angles. How to enforce computational structure?
3. MIN-QAOA even more restricted than MIN-VQA — permits only two Hamiltonians, one of which also acts as observable?

Challenge 1: Hardness of approximation

Quantum Monotone Minimum Satisfying Assignment (QMSA)

Given quantum circuit V accepting non-empty monotone set $S \subseteq \{0,1\}^n$, weight thresholds $g \leq g'$, output:

- YES if $\exists\ x \in \{0,1\}^n$ of Hamming weight at most g accepted by V.
- NO if $\forall\ x \in \{0,1\}^n$ of Hamming weight at most g' are rejected by V.

Previously known:

- $\forall \epsilon > 0$, QMSA is QCMA-hard to approximate within ratio $g'/g \in O(N^{1-\epsilon})$ [G, Kempe, 2012]
- Exploits disperser-based NP-hardness of approximation framework of [Umans 1999] for Σ_2^p

Challenge 1: Hardness of approximation

Quantum Monotone Minimum Satisfying Assignment (QMSA)

Given quantum circuit V accepting non-empty monotone set $S \subseteq \{0,1\}^n$, weight thresholds $g \leq g'$, output:

- YES if $\exists\ x \in \{0,1\}^n$ of Hamming weight at most g accepted by V.
- NO if $\forall\ x \in \{0,1\}^n$ of Hamming weight at most g' are rejected by V.

Previously known:

- $\forall \epsilon > 0$, QMSA is QCMA-hard to approximate within ratio $g'/g \in O(N^{1-\epsilon})$ [G, Kempe, 2012]
- Exploits disperser-based NP-hardness of approximation framework of [Umans 1999] for Σ_2^p

To overcome Challenge 1:

- Reduce QMSA to MIN-VQA via poly-time, many-one reduction
- Maintaining $N^{1-\epsilon}$ hardness ratio will require special attention

Challenge 2: Enforcing computational structure

Revised Goal: Map given QMSA instance (V, g, g') to instance (H, d, d') of MIN-VQA s.t. $\frac{d'}{d} \geq N^{1-\epsilon}$, and

$\exists$ proof y of Hamming weight $\leq g$ accepted by V $\implies$ $\leq d$ VQA levels to get "good" measurement result

$\forall$ proofs y of Hamming weight $\leq g'$, V rejects $\implies$ $> d'$ VQA levels to get "good" measurement result

Challenge 2: Enforcing computational structure

Revised Goal: Map given QMSA instance (V, g, g') to instance (H, d, d') of MIN-VQA s.t. $\frac{d'}{d} \geq N^{1-\epsilon}$, and

$\exists$ proof y of Hamming weight $\leq g$ accepted by V $\implies$ $\leq d$ VQA levels to get "good" measurement result

$\forall$ proofs y of Hamming weight $\leq g'$, V rejects $\implies$ $> d'$ VQA levels to get "good" measurement result

Idea

- Use "hybrid Cook-Levin + Kitaev" circuit-to-Hamiltonian construction

Challenge 2: Enforcing computational structure

Revised Goal: Map given QMSA instance (V, g, g') to instance (H, d, d') of MIN-VQA s.t. $\frac{d'}{d} \geq N^{1-\epsilon}$, and

$\exists$ proof y of Hamming weight $\leq g$ accepted by V $\implies$ $\leq d$ VQA levels to get "good" measurement result

$\forall$ proofs y of Hamming weight $\leq g'$, V rejects $\implies$ $> d'$ VQA levels to get "good" measurement result

Idea

- Use "hybrid Cook-Levin + Kitaev" circuit-to-Hamiltonian construction
- Build set of VQA Hamiltonians $H = P \cup Q \cup F \cup G$, such that for an honest prover:
  1. (Proof) Hamiltonians from P used to prepare proof y

Challenge 2: Enforcing computational structure

Revised Goal: Map given QMSA instance (V, g, g') to instance (H, d, d') of MIN-VQA s.t. $\frac{d'}{d} \geq N^{1-\epsilon}$, and

$\exists$ proof y of Hamming weight $\leq g$ accepted by V $\implies$ $\leq d$ VQA levels to get "good" measurement result

$\forall$ proofs y of Hamming weight $\leq g'$, V rejects $\implies$ $> d'$ VQA levels to get "good" measurement result

Idea

- Use "hybrid Cook-Levin + Kitaev" circuit-to-Hamiltonian construction
- Build set of VQA Hamiltonians $H = P \cup Q \cup F \cup G$, such that for an honest prover:
  1. (Proof) Hamiltonians from P used to prepare proof y
  2. (Quantum verifier V) Hamiltonians from Q used to simulate QMSA verifier V's gates

Challenge 2: Enforcing computational structure

Revised Goal: Map given QMSA instance (V, g, g') to instance (H, d, d') of MIN-VQA s.t. $\frac{d'}{d} \geq N^{1-\epsilon}$, and

$\exists$ proof y of Hamming weight $\leq g$ accepted by V $\implies$ $\leq d$ VQA levels to get "good" measurement result

$\forall$ proofs y of Hamming weight $\leq g'$, V rejects $\implies$ $> d'$ VQA levels to get "good" measurement result

Idea

- Use "hybrid Cook-Levin + Kitaev" circuit-to-Hamiltonian construction
- Build set of VQA Hamiltonians $H = P \cup Q \cup F \cup G$, such that for an honest prover:
  1. (Proof) Hamiltonians from P used to prepare proof y
  2. (Quantum verifier V) Hamiltonians from Q used to simulate QMSA verifier V's gates
  3. (2D clock) Hamiltonians from $F \cup G$ implement "2D clock" to track time and preserve hardness gap d'/d

Challenge 2: Enforcing computational structure

VQA Hamiltonians act on four registers: A (proof), B (workspace), C (clock 1) and D (clock 2)

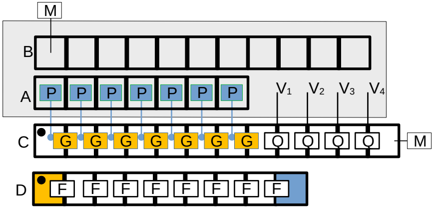

Challenge 2: Enforcing computational structure

VQA Hamiltonians act on four registers: A (proof), B (workspace), C (clock 1) and D (clock 2)

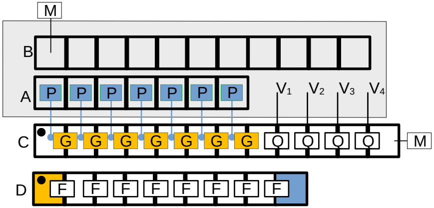

- Set of "proof" Hamiltonians, P, consists of (e.g.)

$X_{A_j} \otimes |1\rangle\langle 1|_{C_j} \otimes |1\rangle\langle 1|_{D_{|D|}}$ $\xRightarrow{\text{evolve for } \theta=\pi/2}$ apply X to jth proof qubit if clocks C and D are j and $|D|$, resp.

Challenge 2: Enforcing computational structure

VQA Hamiltonians act on four registers: A (proof), B (workspace), C (clock 1) and D (clock 2)

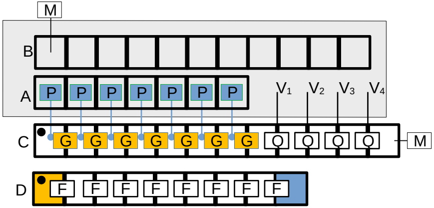

- Set of "proof" Hamiltonians, P, consists of (e.g.)

$$X_{A_j} \otimes |1\rangle\langle 1|_{C_j} \otimes |1\rangle\langle 1|_{D_{|D|}} \xRightarrow{\text{evolve for } \theta=\pi/2}$$ apply X to jth proof qubit if clocks C and D are j and $|D|$, resp.

- Set of "quantum verifier" Hamiltonians, Q, consists of (e.g.)

$$(V_j)_{AB} \otimes |01\rangle\langle 10|_{C_{|A|+j,|A|+j+1}} + (V_j^\dagger)_{AB} \otimes |10\rangle\langle 01|_{C_{|A|+j,|A|+j+1}} \xRightarrow{\text{evolve for } \theta=\pi/2}$$

apply jth gate of verifier V, update clock C from $|A|+j$ to $|A|+j+1$

Challenge 2: Enforcing computational structure

VQA Hamiltonians act on four registers: A (proof), B (workspace), C (clock 1) and D (clock 2)

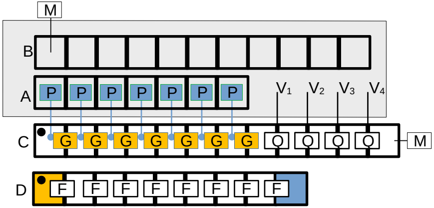

- Set of "proof" Hamiltonians, P, consists of (e.g.)

$$X_{A_j} \otimes |1\rangle\langle 1|_{C_j} \otimes |1\rangle\langle 1|_{D_{|D|}} \xRightarrow{\text{evolve for } \theta=\pi/2}$$ apply X to jth proof qubit if clocks C and D are j and $|D|$, resp.

- Set of "quantum verifier" Hamiltonians, Q, consists of (e.g.)

$$(V_j)_{AB} \otimes |01\rangle\langle 10|_{C_{|A|+j,|A|+j+1}} + (V_j^\dagger)_{AB} \otimes |10\rangle\langle 01|_{C_{|A|+j,|A|+j+1}} \xRightarrow{\text{evolve for } \theta=\pi/2}$$

apply jth gate of verifier V, update clock C from $|A|+j$ to $|A|+j+1$

- Observable M measures output qubit of V when clock C set to $|C|$

Honest provers actions, given instance (V, g, g') of QMSA

1. Prepare proof y by flipping each appropriate bit of register A
   - Takes HammingWeight(y) many Hamiltonian evolutions from set P
2. Simulate each gate of verifier $V = V_L \cdots V_1$
   - Takes L many Hamiltonian evolutions from set Q
3. Observable M now applies energy penalty if V would reject y

Honest provers actions, given instance (V, g, g') of QMSA

1. Prepare proof y by flipping each appropriate bit of register A
   - Takes HammingWeight(y) many Hamiltonian evolutions from set P
2. Simulate each gate of verifier $V = V_L \cdots V_1$
   - Takes L many Hamiltonian evolutions from set Q
3. Observable M now applies energy penalty if V would reject y

Bad news: Honest prover above applies HammingWeight(y)$+L$ evolutions, so ratio obtained scales as

$$\frac{g' + L}{g + L} \rightarrow 1 \text{ if } L \in \omega(g).$$

Honest provers actions, given instance (V, g, g') of QMSA

1. Prepare proof y by flipping each appropriate bit of register A
   - Takes HammingWeight(y) many Hamiltonian evolutions from set P
2. Simulate each gate of verifier $V = V_L \cdots V_1$
   - Takes L many Hamiltonian evolutions from set Q
3. Observable M now applies energy penalty if V would reject y

Bad news: Honest prover above applies HammingWeight(y)$+L$ evolutions, so ratio obtained scales as

$$\frac{g' + L}{g + L} \to 1 \text{ if } L \in \omega(g).$$

Fix: Use 2D clock to make flipping each bit of proof "more costly" without blowing up encoding size:

$$\text{New hardness ratio: } \frac{g' \,|D| + L}{g \,|D| + L} \approx \frac{g'}{g} \text{ for } |D| \in \omega(L), \quad (1)$$

Soundness for dishonest prover

Computation Subspace Preservation Lemma

For any sequence of angles $\theta_j \in \mathbb{R}$ and Hamiltonians $H_j \in P \cup Q \cup F \cup G$,

$$e^{i\theta_m H_m} \cdots e^{i\theta_2 H_2} e^{i\theta_1 H_1} |0 \cdots 0\rangle_{ABCD}$$

is in span of states from

$$S := \left\{ V_{s-|A|} \cdots V_1 |y\rangle_A |0 \cdots 0\rangle_B |\widetilde{s}\rangle_C |\widetilde{t}\rangle_D \mid y \in \{0,1\}^{|A|}, s \in \{1, \ldots, |C|\}, t \in \{1, \ldots, |D|\} \right\}. \quad (2)$$

In words:

- Any sequence of Hamiltonian evolutions keeps us in "logical computation space" S.
- Implication: Forces prover to essentially follow honest strategy

Challenge 3: Extending to QAOA

For QAOA:

- Only 2 Hamiltonians allowed, H_b (driving Hamiltonian) and H_c (cost Hamiltonian),
- start state *implicitly* given as unique ground state of H_b
- no separate observable M.

Challenge 3: Extending to QAOA

For QAOA:

- Only 2 Hamiltonians allowed, H_b (driving Hamiltonian) and H_c (cost Hamiltonian),
- start state *implicitly* given as unique ground state of H_b
- no separate observable M.

Core idea: Alternate even/odd steps of honest prover's actions, i.e. H_b does even steps, H_c odd steps.

Challenge 3: Extending to QAOA

For QAOA:

- Only 2 Hamiltonians allowed, H_b (driving Hamiltonian) and H_c (cost Hamiltonian),
- start state *implicitly* given as unique ground state of H_b
- no separate observable M.

Core idea: Alternate even/odd steps of honest prover's actions, i.e. H_b does even steps, H_c odd steps.

Under the hood, build on MIN-VQA construction as follows:

1. Make all odd (respectively, even) local terms H_i pairwise commute.
2. Introduce 3-cyclic local terms G_j which encode *multiple* logical actions (instead of just 2)
3. Add constraints to H_b to ensure its unique ground state is correct start state.
4. M added as local term to H_c, but scaled larger than all other terms in H_c.

Summary

- Estimating the optimal depth of a VQA/QAOA ansatz is intractable, even with large multiplicative error
- Formally, QCMA-hard within multiplicative error $N^{1-\epsilon}$ for any $\epsilon > 0$.
- First natural hardness of approximation results for QCMA

Summary

- Estimating the optimal depth of a VQA/QAOA ansatz is intractable, even with large multiplicative error
- Formally, QCMA-hard within multiplicative error $N^{1-\epsilon}$ for any $\epsilon > 0$.
- First natural hardness of approximation results for QCMA

Open questions

- NP-hardness of approximation for QAOA depth for *classical* cost Hamiltonian?
- Good *heuristics* for approximating depth in practice?
- How hard is optimal depth approximation in *noisy* setting?
- Hardness of approximation for other QCMA-complete problems?

Summary

- Estimating the optimal depth of a VQA/QAOA ansatz is intractable, even with large multiplicative error
- Formally, QCMA-hard within multiplicative error $N^{1-\epsilon}$ for any $\epsilon > 0$.
- First natural hardness of approximation results for QCMA

Open questions

- NP-hardness of approximation for QAOA depth for *classical* cost Hamiltonian?
- Good *heuristics* for approximating depth in practice?
- How hard is optimal depth approximation in *noisy* setting?
- Hardness of approximation for other QCMA-complete problems?

"Moral" questions

- Obtained hardness of approximation without quantum PCP. In "classical SAT" language:
  - Leveraged hardness of approximation relative to Hamming weight of satisfying assignments
  - In contrast, "classic PCP for SAT" gives hardness of approximation relative to # clauses satisfied
- Quantum complexity theory — hero or villain?