15249C

Sandia National Laboratories

The Quantum and Classical Streaming Complexity of Quantum and Classical Max-Cut

Presented by:

John Kallaugher[1] Ojas Parekh[1]

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell

dia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly ow subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Streaming Graph Algorithms

Algorithms for graphs received *one edge at a time*.

Streaming Graph Algorithms

Algorithms for graphs received *one edge at a time*.

Streaming Graph Algorithms

Algorithms for graphs received *one edge at a time*.

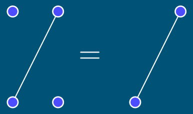

Streaming Graph Algorithms

Algorithms for graphs received *one edge at a time*.

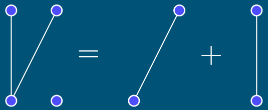

Streaming Graph Algorithms

Algorithms for graphs received *one edge at a time*.

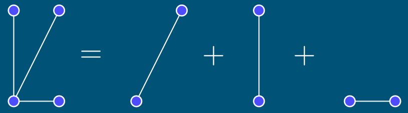

Streaming Graph Algorithms

Algorithms for graphs received *one edge at a time*.

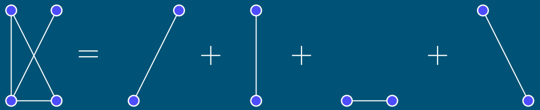

Streaming Graph Algorithms

Algorithms for graphs received *one edge at a time*.

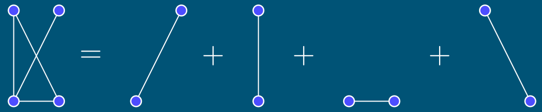

Theoretical study is mostly concerned with *space complexity*.

Max-Cut

$G = (V, E)$, find the largest number of edges crossing any partition $V = V_1 \sqcup V_2$.

Max-Cut

$G = (V, E)$, find the largest number of edges crossing any partition $V = V_1 \sqcup V_2$.

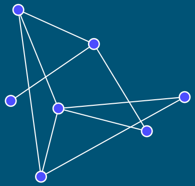

Max-Cut

$G = (V, E)$, find the largest number of edges crossing any partition $V = V_1 \sqcup V_2$.

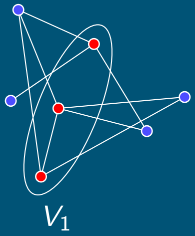

Max-Cut

$G = (V, E)$, find the largest number of edges crossing any partition $V = V_1 \sqcup V_2$.

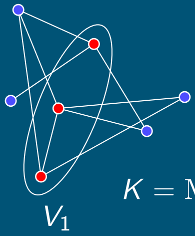

$K = \textsc{Max-Cut}(G) = 8$

Max-Cut

$G = (V, E)$, find the largest number of edges crossing any partition $V = V_1 \sqcup V_2$.

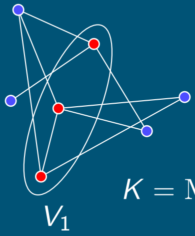

$K = \text{Max-Cut}(G) = 8$

- A canonical example of a *constraint satisfaction problem* (CSP).

Max-Cut

$G = (V, E)$, find the largest number of edges crossing any partition $V = V_1 \sqcup V_2$.

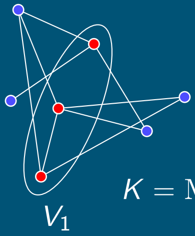

$K = \text{Max-Cut}(G) = 8$

- A canonical example of a *constraint satisfaction problem* (CSP).
- NP-hard to approximate.

Max-Cut

$G = (V, E)$, find the largest number of edges crossing any partition $V = V_1 \sqcup V_2$.

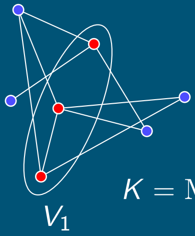

$K = \textsc{Max-Cut}(G) = 8$

- A canonical example of a *constraint satisfaction problem* (CSP).
- NP-hard to approximate.
- How much space does a streaming algorithm need to return $K' \in (K/\gamma, K)$?

Streaming Complexity of MAX-CUT

- 2-approximation is trivial in $O(\log n)$ space.

Streaming Complexity of Max-Cut

- 2-approximation is trivial in $O(\log n)$ space.
- $(1+\varepsilon)$-approximation in $\widetilde{O}(n)$ space with sparsifiers.

Streaming Complexity of Max-Cut

- 2-approximation is trivial in $O(\log n)$ space.
- $(1+\varepsilon)$-approximation in $\widetilde{O}(n)$ space with sparsifiers.
- Gap gradually closed in [Kogan, Krauthgamer '15], [Kapralov, Khanna, Sudan '15], [Kapralov, Khanna, Sudan, Velingker '17].

Streaming Complexity of Max-Cut

- 2-approximation is trivial in $O(\log n)$ space.
- $(1+\varepsilon)$-approximation in $\widetilde{O}(n)$ space with sparsifiers.
- Gap gradually closed in [Kogan, Krauthgamer '15], [Kapralov, Khanna, Sudan '15], [Kapralov, Khanna, Sudan, Velingker '17].
- Resolved in [Kapralov, Krachun '19]. $(2-\varepsilon)$-approximation requires $\Omega(n)$ space.

Our Results

We generalize these results in two directions.

	Max-Cut		
Approximation Factor	$2+\varepsilon$	$2-\varepsilon$	
Classical Algorithm	$O(\log n)$	$\Omega(n)$	

Our Results

We generalize these results in two directions.

	Max-Cut		
Approximation Factor	$2+\varepsilon$	$2-\varepsilon$	
Classical Algorithm	$O(\log n)$	$\Omega(n)$	
Quantum Algorithm	$O(\log n)$	$\Omega(n)$	

- Lower bounds for *quantum* streaming algorithms.

Our Results

We generalize these results in two directions.

	MAX-CUT		QUANTUM MAX-CUT
Approximation Factor	$2+\varepsilon$	$2-\varepsilon$	$2-\varepsilon$
Classical Algorithm	$O(\log n)$	$\Omega(n)$	$\Omega(n)$
Quantum Algorithm	$O(\log n)$	$\Omega(n)$	$\Omega(n)$

- Lower bounds for *quantum* streaming algorithms.
  - Key tool: managing the evolution of the Fourier coefficients of a quantum protocol under the application of a quantum channel.

Our Results

We generalize these results in two directions.

	Max-Cut		Quantum Max-Cut	
Approximation Factor	$2+\varepsilon$	$2-\varepsilon$	$2+\varepsilon$	$2-\varepsilon$
Classical Algorithm	$O(\log n)$	$\Omega(n)$	$O(\log n)$	$\Omega(n)$
Quantum Algorithm	$O(\log n)$	$\Omega(n)$	$O(\log n)$	$\Omega(n)$

- Lower bounds for *quantum* streaming algorithms.
    - Key tool: managing the evolution of the Fourier coefficients of a quantum protocol under the application of a quantum channel.
- And for Quantum Max-Cut.

Our Results

We generalize these results in two directions.

	Max-Cut		Quantum Max-Cut	
Approximation Factor	$2+\varepsilon$	$2-\varepsilon$	$2+\varepsilon$	$2-\varepsilon$
Classical Algorithm	$O(\log n)$	$\Omega(n)$	$O(\log n)$	$\Omega(n)$
Quantum Algorithm	$O(\log n)$	$\Omega(n)$	$O(\log n)$	$\Omega(n)$

- Lower bounds for *quantum* streaming algorithms.
  - Key tool: managing the evolution of the Fourier coefficients of a quantum protocol under the application of a quantum channel.
- And for Quantum Max-Cut.
  - Matching upper bound for $(2+\varepsilon)$-approximation.

Quantum Streaming Algorithms

Streaming algorithms that maintain *quantum* state.

Quantum Streaming Algorithms

Streaming algorithms that maintain *quantum* state.

$$|0\rangle^{\otimes k} = \phi$$

Quantum Streaming Algorithms

Streaming algorithms that maintain *quantum* state.

$$|0\rangle^{\otimes k} = \phi \underset{u_1}{\Rightarrow} \mathcal{N}_{u_1}(\phi)$$

Quantum Streaming Algorithms

Streaming algorithms that maintain *quantum* state.

$$|0\rangle^{\otimes k} = \phi \underset{u_1}{\Rightarrow} \mathcal{N}_{u_1}(\phi) \underset{u_2}{\Rightarrow} \mathcal{N}_{u_2} \circ \mathcal{N}_{u_1}(\phi)$$

Quantum Streaming Algorithms

Streaming algorithms that maintain *quantum* state.

$$|0\rangle^{\otimes k} = \phi \underset{u_1}{\Rightarrow} \mathcal{N}_{u_1}(\phi) \underset{u_2}{\Rightarrow} \mathcal{N}_{u_2} \circ \mathcal{N}_{u_1}(\phi)$$

- [Le Gall '06], [Gavinsky, Kempe, Kerenidis, Raz, de Wolf '08] Exponential space advantage possible.

Quantum Streaming Algorithms

Streaming algorithms that maintain *quantum* state.

$$|0\rangle^{\otimes k} = \phi \underset{u_1}{\Rightarrow} \mathcal{N}_{u_1}(\phi) \underset{u_2}{\Rightarrow} \mathcal{N}_{u_2} \circ \mathcal{N}_{u_1}(\phi)$$

- [Le Gall '06], [Gavinsky, Kempe, Kerenidis, Raz, de Wolf '08] Exponential space advantage possible.
- [Montanaro '16], [Hamoudi, Magniez '19], [**K.** '22] Polynomial advantages known for "natural" problems.

Distributed Implicit Hidden Partition (DIHP)

Introduced in [Kapralov, Krachun '19].

Distributed Implicit Hidden Partition (DIHP)

Introduced in [Kapralov, Krachun '19].

- T-player sequential communication game. Player t gets a random partial matching M_t on αn vertices, $|V| = n$, $\alpha = \Theta(1/\sqrt{\alpha})$.

Distributed Implicit Hidden Partition (DIHP)

Introduced in [Kapralov, Krachun '19].

YES:

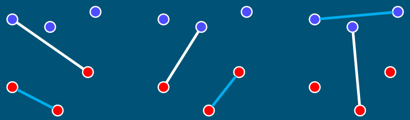

- T-player sequential communication game. Player t gets a random partial matching M_t on αn vertices, $|V| = n$, $\alpha = \Theta(1/\sqrt{\alpha})$.
- **YES** case: there is a secret string $x \in \{0,1\}^V$. Player t gets $y = M_t x$, i.e. $y_{uv} = x_u \oplus x_v$ for each edge $uv \in M_t$.

Distributed Implicit Hidden Partition (DIHP)

Introduced in [Kapralov, Krachun '19].

YES:

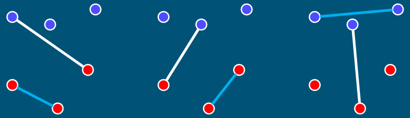

NO:

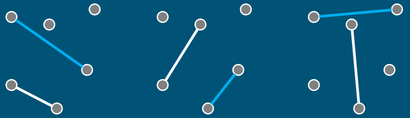

- T-player sequential communication game. Player t gets a random partial matching M_t on αn vertices, $|V| = n$, $\alpha = \Theta(1/\sqrt{\alpha})$.
- **YES** case: there is a secret string $x \in \{0,1\}^V$. Player t gets $y = M_t x$, i.e. $y_{uv} = x_u \oplus x_v$ for each edge $uv \in M_t$.
- **NO** case: Player t gets $y \in \{0,1\}^{M_t}$ at random.

Reduction from MAX-CUT

Consider the graph of the m edges labelled 1.

YES: **NO**:

Reduction from MAX-CUT

Consider the graph of the m edges labelled 1.

- **YES** case: x gives a perfect cut of this graph.

Reduction from MAX-CUT

Consider the graph of the m edges labelled 1.

- **YES** case: x gives a perfect cut of this graph.
- **NO** case: Graph is close to random, best cut is $m(1/2 + \Theta\left(1/\sqrt{T}\right))$.

Reduction from Max-Cut

Consider the graph of the m edges labelled 1.

- **YES** case: x gives a perfect cut of this graph.
- **NO** case: Graph is close to random, best cut is $m(1/2 + \Theta\left(1/\sqrt{T}\right))$.
- DIHP reduces to $(2 - \Theta\left(1/\sqrt{T}\right))$ approximating Max-Cut.

Fourier Coefficients

[Kapralov, Krachun '19] Solving DIHP requires learning many even parities of the underlying partition (if there is one).

Fourier Coefficients

[Kapralov, Krachun '19] Solving DIHP requires learning many even parities of the underlying partition (if there is one).

- Intuition: in a **YES** case, each edge label given to player t is "consistent" with the corresponding vertex labels.

Fourier Coefficients

[Kapralov, Krachun '19] Solving DIHP requires learning many even parities of the underlying partition (if there is one).

- Intuition: in a **YES** case, each edge label given to player t is "consistent" with the corresponding vertex labels.
- The "knowledge" player t has about these parities can be quantified in terms of *Fourier* coefficients.

Fourier Analysis (of Quantum Protocols)

Given a protocol, we can write down a function $f_t : \{0,1\}^n \to \mathbb{C}^{\beta \times \beta}$ that gives the density matrix $f_t(x)$ of player t's message if the problem is in a **YES** state with vertex labels x.

Fourier Analysis (of Quantum Protocols)

Given a protocol, we can write down a function $f_t : \{0,1\}^n \to \mathbb{C}^{\beta \times \beta}$ that gives the density matrix $f_t(x)$ of player t's message if the problem is in a **YES** state with vertex labels x.

$f_{t-1}(x) \Rightarrow$ 

$\Rightarrow f_t(x)$

$$\Uparrow$$

$$M_t x$$

Fourier Analysis (of Quantum Protocols)

Given a protocol, we can write down a function $f_t : \{0,1\}^n \to \mathbb{C}^{\beta \times \beta}$ that gives the density matrix $f_t(x)$ of player t's message if the problem is in a **YES** state with vertex labels x.

$$f_{t-1}(x) \Rightarrow \boxed{\mathcal{Q}} \Rightarrow f_t(x)$$

$$\Uparrow$$

$$M_t x$$

For each set S of vertices we consider the *Fourier coefficient* $\widehat{f}_t(S) = \mathbb{E}_x \left[(-1)^{S \cdot x} f_t(x)\right]$.

Hypercontractivity

Suppose player t ignored everything from player $t-1$. Then we would have

$$\sum_{|S|=k} \|\widehat{f}_t(S)\|_1^2 \leq \binom{\beta}{k}$$

for β-qubit messages by *hypercontractivity*.

Hypercontractivity

Suppose player t ignored everything from player $t-1$. Then we would have

$$\sum_{|S|=k} \|\widehat{f}_t(S)\|_1^2 \leq \binom{\beta}{k}$$

for β-qubit messages by *hypercontractivity*.

- This would suffice to make player $t+1$ learn very little if $\beta \ll n$.

Hypercontractivity

Suppose player t ignored everything from player $t-1$. Then we would have

$$\sum_{|S|=k} \|\widehat{f_t}(S)\|_1^2 \leq \binom{\beta}{k}$$

for β-qubit messages by *hypercontractivity*.

- This would suffice to make player $t+1$ learn very little if $\beta \ll n$.
- But we need to control how player t incorporates a message from player t.

Evolution of Fourier Coefficients

A key ingredient of the classical proof is the fact that, because the input distribution is known, the protocol can be assumed to be *deterministic*.

Evolution of Fourier Coefficients

A key ingredient of the classical proof is the fact that, because the input distribution is known, the protocol can be assumed to be *deterministic*.

$$f_{t-1}(x) \Rightarrow$$

$$\Rightarrow a(f_{t-1}(x), M_t x)$$

$$\Uparrow$$

$$M_t x$$

Evolution of Fourier Coefficients

A key ingredient of the classical proof is the fact that, because the input distribution is known, the protocol can be assumed to be *deterministic*.

$$f_{t-1}(x) \Rightarrow$$

$$\Rightarrow a(f_{t-1}(x), M_t x)$$

$$\Uparrow$$

$$M_t x$$

In the quantum setting the possibility of e.g. measurements means that we have to consider player t applying an arbitrary *quantum channel*.

Quantum Channels

A quantum channel represents any realizable transformation of a density matrix on β qubits.

$$f_{t-1}(x) \Rightarrow$$

$$\Rightarrow \mathcal{A}^{t}_{M_t x}(f_{t-1}(x))$$

$$\Uparrow$$

$$M_t x$$

Quantum Channels

A quantum channel represents any realizable transformation of a density matrix on β qubits.

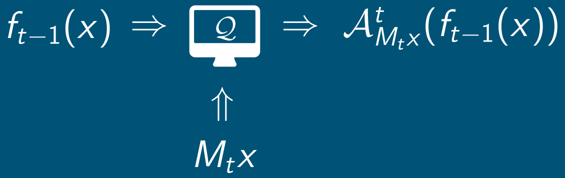

Our idea is to apply Fourier analysis to the family of channels each player might apply.

Fourier Transform of a Quantum Channel

Let $(\mathcal{N}_y)_{y\in\{0,1\}^n}$ be a family of channels. We extend the Boolean Fourier transform by defining:

$$\widehat{\mathcal{N}}_S = \mathop{\mathbb{E}}_{y}\left[(-1)^{S\cdot y}\mathcal{N}_y\right]$$

Fourier Transform of a Quantum Channel

Let $(\mathcal{N}_y)_{y\in\{0,1\}^n}$ be a family of channels. We extend the Boolean Fourier transform by defining:

$$\widehat{\mathcal{N}}_S = \mathop{\mathbb{E}}_{y}\left[(-1)^{S\cdot y}\mathcal{N}_y\right]$$

A convolution theorem (analogous to one for products of scalar-valued functions) applies:

$$\widehat{\mathcal{N}f}(S) = \sum_{U\in\{0,1\}^n} \widehat{\mathcal{N}}_U \widehat{f}(U\oplus S)$$

Fourier Coefficients Under Quantum Evolution

$\mathbb{E}_y\left[(-1)^{S\cdot y} f_t(x)\right]$ is non-zero iff U corresponds to a set of edges in M_t.

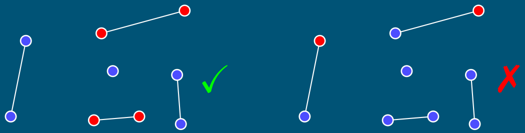

Fourier Coefficients Under Quantum Evolution

$\mathbb{E}_y\left[(-1)^{S\cdot y} f_t(x)\right]$ is non-zero iff U corresponds to a set of edges in M_t.

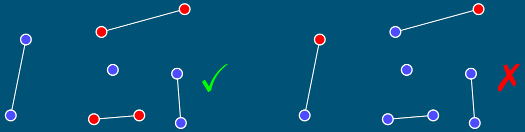

So each Fourier coefficient $\widehat{f}_{t-1}(S)$ can be seen as generating one Fourier coefficient $\widehat{f}_t(S \oplus M_t^{\mathrm{tr}} y)$ for each subset y of M_t.

DIHP Lower Bound

Combinatorics and matrix hypercontractivity gives us that $\|f_{t-1}(S)\|_1$ will be small for all S matched by player t's edges if $\#$ of qubits $\beta \ll n/2^{1/T}$.

DIHP Lower Bound

Combinatorics and matrix hypercontractivity gives us that $\|f_{t-1}(S)\|_1$ will be small for all S matched by player t's edges if # of qubits $\beta \ll n/2^{1/T}$.

- DIHP requires $n/2^{1/T}$ (qu)bits of quantum *or* classical communication.

DIHP Lower Bound

Combinatorics and matrix hypercontractivity gives us that $\|f_{t-1}(S)\|_1$ will be small for all S matched by player t's edges if # of qubits $\beta \ll n/2^{1/T}$.

- DIHP requires $n/2^{1/T}$ (qu)bits of quantum *or* classical communication.
- So $(2-\varepsilon)$-approximation of Max-Cut requires $n/2^{O(\varepsilon)}$ space.

Quantum Generalizations of CSPs

The quantum generalization of a constraint satisfaction problem is a *local Hamiltonian* problem.

Quantum Generalizations of CSPs

The quantum generalization of a constraint satisfaction problem is a *local Hamiltonian* problem.

$$H = \sum I \otimes L$$

- A sum of "local" terms operating on a small number of qubits each.

Quantum Generalizations of CSPs

The quantum generalization of a constraint satisfaction problem is a *local Hamiltonian* problem.

$$H = \sum I \otimes L$$

- A sum of "local" terms operating on a small number of qubits each.
- Objective is to find minimum/maximum $x^* H x$ among n-qubit states x.

Quantum Generalizations of CSPs

The quantum generalization of a constraint satisfaction problem is a *local Hamiltonian* problem.

$$H = \sum I \otimes L$$

- A sum of "local" terms operating on a small number of qubits each.
- Objective is to find minimum/maximum $x^* H x$ among n-qubit states x.
- Many important problems in physics take this form.

Quantum Max-Cut

One of these is the *anti-ferromagnetic Heisenberg model*.

Quantum Max-Cut

One of these is the *anti-ferromagnetic Heisenberg model*.

- Searching for the minimum energy of this creates a Hamiltonian maximization problem that returns half the cut value of any *classical* string viewed as a partition.

Quantum Max-Cut

One of these is the *anti-ferromagnetic Heisenberg model*.

- Searching for the minimum energy of this creates a Hamiltonian maximization problem that returns half the cut value of any *classical* string viewed as a partition.
- Also referred to as Quantum Max-Cut.

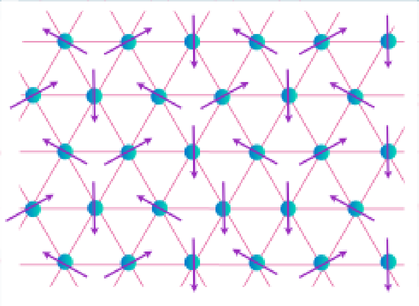

Figure: Anti-ferromagnetic Heisenberg model: roughly neighboring quantum particles aim to align in opposite directions.
[Image: Sachdev, arXiv:1203.4565]

Reducing DIHP to Quantum Max-Cut

The Quantum Max-Cut value of a graph is always at least 1/2 its Max-Cut value, so a **YES** instance of DIHP has value $m/4$.

Reducing DIHP to Quantum Max-Cut

The Quantum Max-Cut value of a graph is always at least 1/2 its Max-Cut value, so a **YES** instance of DIHP has value $m/4$. For the **NO** case we will need the following SDP:

$$\max_{f:V\to S^{n-1}} \sum_{uv\in E} -\langle f(u), f(v)\rangle$$

Reducing DIHP to Quantum Max-Cut

The Quantum Max-Cut value of a graph is always at least 1/2 its Max-Cut value, so a **YES** instance of DIHP has value $m/4$. For the **NO** case we will need the following SDP:

$$\max_{f:V\to S^{n-1}} \sum_{uv\in E} -\langle f(u), f(v)\rangle$$

This is a shifted version of the Goemans-Williamson SDP for Max-Cut. We use two properties:

Reducing DIHP to Quantum Max-Cut

The Quantum Max-Cut value of a graph is always at least 1/2 its Max-Cut value, so a **YES** instance of DIHP has value $m/4$. For the **NO** case we will need the following SDP:

$$\max_{f:V\to S^{n-1}} \sum_{uv\in E} -\langle f(u), f(v)\rangle$$

This is a shifted version of the Goemans-Williamson SDP for Max-Cut. We use two properties:

- If it is ε, Quantum Max-Cut value is $m/2 + O(\varepsilon)$.

Reducing DIHP to Quantum Max-Cut

The Quantum Max-Cut value of a graph is always at least 1/2 its Max-Cut value, so a **YES** instance of DIHP has value $m/4$. For the **NO** case we will need the following SDP:

$$\max_{f:V\to S^{n-1}} \sum_{uv\in E} -\langle f(u), f(v)\rangle$$

This is a shifted version of the Goemans-Williamson SDP for Max-Cut. We use two properties:

- If it is ε, Quantum Max-Cut value is $m/2 + O(\varepsilon)$.
- If Max-Cut value is $m/2 + O(\varepsilon)$, SDP is $O(\varepsilon)$.

Reducing DIHP to Quantum Max-Cut

The Quantum Max-Cut value of a graph is always at least 1/2 its Max-Cut value, so a **YES** instance of DIHP has value $m/4$. For the **NO** case we will need the following SDP:

$$\max_{f:V\to S^{n-1}} \sum_{uv\in E} -\langle f(u), f(v)\rangle$$

This is a shifted version of the Goemans-Williamson SDP for Max-Cut. We use two properties:

- If it is ε, Quantum Max-Cut value is $m/2 + O(\varepsilon)$.
- If Max-Cut value is $m/2 + O(\varepsilon)$, SDP is $O(\varepsilon)$.

So **NO** instances have Quantum Max-Cut value at most $m(1/2 + \Theta\left(1/\sqrt{T}\right))$ too!

Conclusion

	Max-Cut		Quantum Max-Cut	
Approximation Factor	$2+\varepsilon$	$2-\varepsilon$	$2+\varepsilon$	$2-\varepsilon$
Classical Algorithm	$O(\log n)$	$\Omega(n)$	$\mathbf{O(\log n)}$	$\mathbf{\Omega(n)}$
Quantum Algorithm	$O(\log n)$	$\mathbf{\Omega(n)}$	$\mathbf{O(\log n)}$	$\mathbf{\Omega(n)}$

Conclusion

	Max-Cut		Quantum Max-Cut	
Approximation Factor	$2+\varepsilon$	$2-\varepsilon$	$2+\varepsilon$	$2-\varepsilon$
Classical Algorithm	$O(\log n)$	$\Omega(n)$	$\mathbf{O(\log n)}$	$\mathbf{\Omega(n)}$
Quantum Algorithm	$O(\log n)$	$\mathbf{\Omega(n)}$	$\mathbf{O(\log n)}$	$\mathbf{\Omega(n)}$

- [Chou, Golovnev, Sudan, Velingker, Velusamy '22] gives linear lower bounds for a larger class of CSPs. Do these results translate?

Conclusion

	MAX-CUT		QUANTUM MAX-CUT	
Approximation Factor	$2+\varepsilon$	$2-\varepsilon$	$2+\varepsilon$	$2-\varepsilon$
Classical Algorithm	$O(\log n)$	$\Omega(n)$	$\mathbf{O(\log n)}$	$\mathbf{\Omega(n)}$
Quantum Algorithm	$O(\log n)$	$\mathbf{\Omega(n)}$	$\mathbf{O(\log n)}$	$\mathbf{\Omega(n)}$

- [Chou, Golovnev, Sudan, Velingker, Velusamy '22] gives linear lower bounds for a larger class of CSPs. Do these results translate?
- Weighted QUANTUM MAX-CUT.

Conclusion

	Max-Cut		Quantum Max-Cut	
Approximation Factor	$2+\varepsilon$	$2-\varepsilon$	$2+\varepsilon$	$2-\varepsilon$
Classical Algorithm	$O(\log n)$	$\Omega(n)$	$\mathbf{O(\log n)}$	$\mathbf{\Omega(n)}$
Quantum Algorithm	$O(\log n)$	$\mathbf{\Omega(n)}$	$\mathbf{O(\log n)}$	$\mathbf{\Omega(n)}$

- [Chou, Golovnev, Sudan, Velingker, Velusamy '22] gives linear lower bounds for a larger class of CSPs. Do these results translate?
- Weighted Quantum Max-Cut.
- Can we sparsify (some) 2-local Hamiltonians?

Conclusion

	Max-Cut		Quantum Max-Cut	
Approximation Factor	$2+\varepsilon$	$2-\varepsilon$	$2+\varepsilon$	$2-\varepsilon$
Classical Algorithm	$O(\log n)$	$\Omega(n)$	**$O(\log n)$**	**$\Omega(n)$**
Quantum Algorithm	$O(\log n)$	**$\Omega(n)$**	**$O(\log n)$**	**$\Omega(n)$**

- [Chou, Golovnev, Sudan, Velingker, Velusamy '22] gives linear lower bounds for a larger class of CSPs. Do these results translate?
- Weighted Quantum Max-Cut.
- Can we sparsify (some) 2-local Hamiltonians?
- If we can, can we solve them in $O(n)$ space?