Quantum Cut Sparsifiers




Arpon Basu
Princeton

Joshua Brakensiek
UC Berkeley

Pravesh Kothari
Princeton

Aaron Putterman
Harvard

What is Sparsification?

- Compressing objects (graphs, hypergraphs, codes, …) while retaining essential properties

- Running example throughout Talk: Graph Sparsification

- Graph $G = G(V, E)$ with vertex set V and edge set E
  - V has n vertices, and the edge set E might be weighted

- For any $S \subset V$, let $\mathrm{cut}_G(S)$ be the total number/weight of edges leaving S

Cuts in Graphs

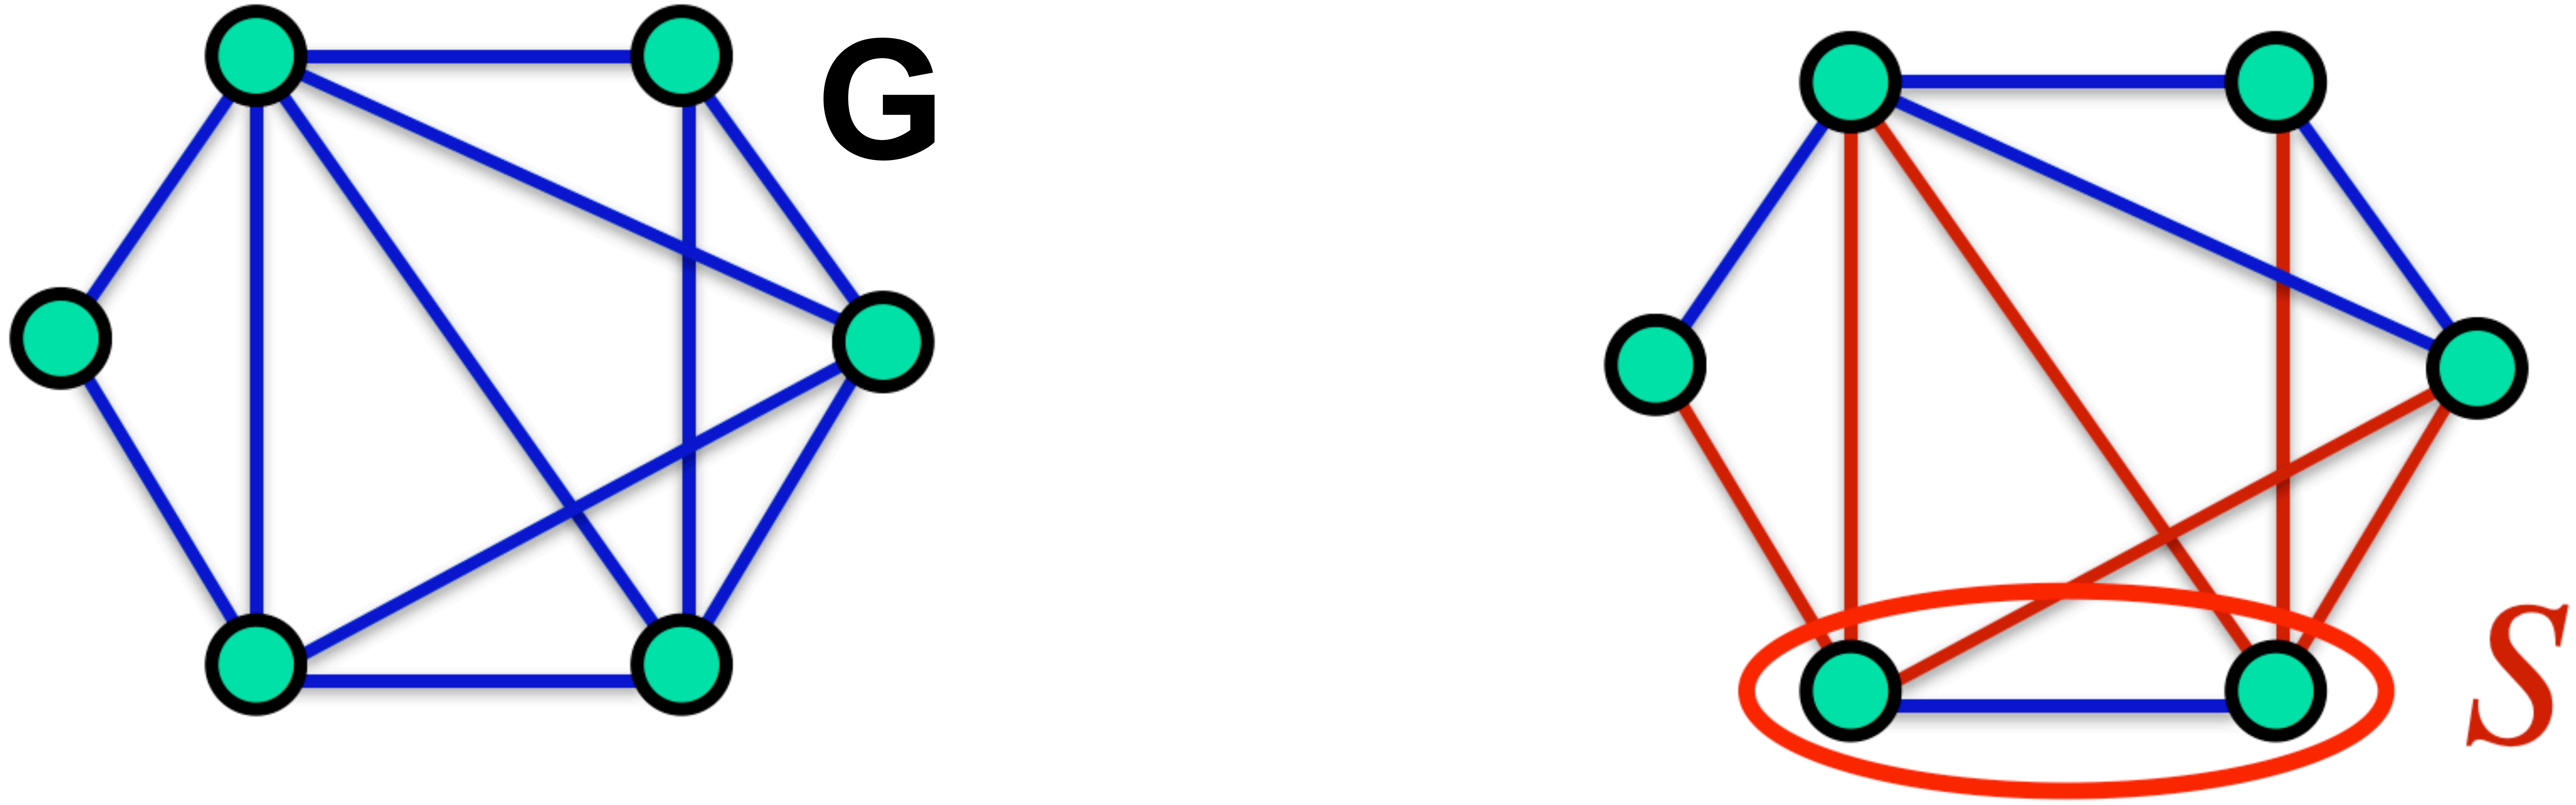

$\mathrm{cut}_G(S)$ equals the sum of the weights of all the red edges

Figure due to Aaron Putterman

Cut Sparsifiers

- For a graph $G = G(V, E)$, a cut sparsifier is a sparse graph $\tilde{G} = \tilde{G}(V, \tilde{E})$ such that for **every** $S \subset V$ we have

$$\mathrm{cut}_{\tilde{G}}(S) \in [0.99, 1.01] \cdot \mathrm{cut}_G(S)$$

Edges in $\tilde{G}$ have higher weights to compensate

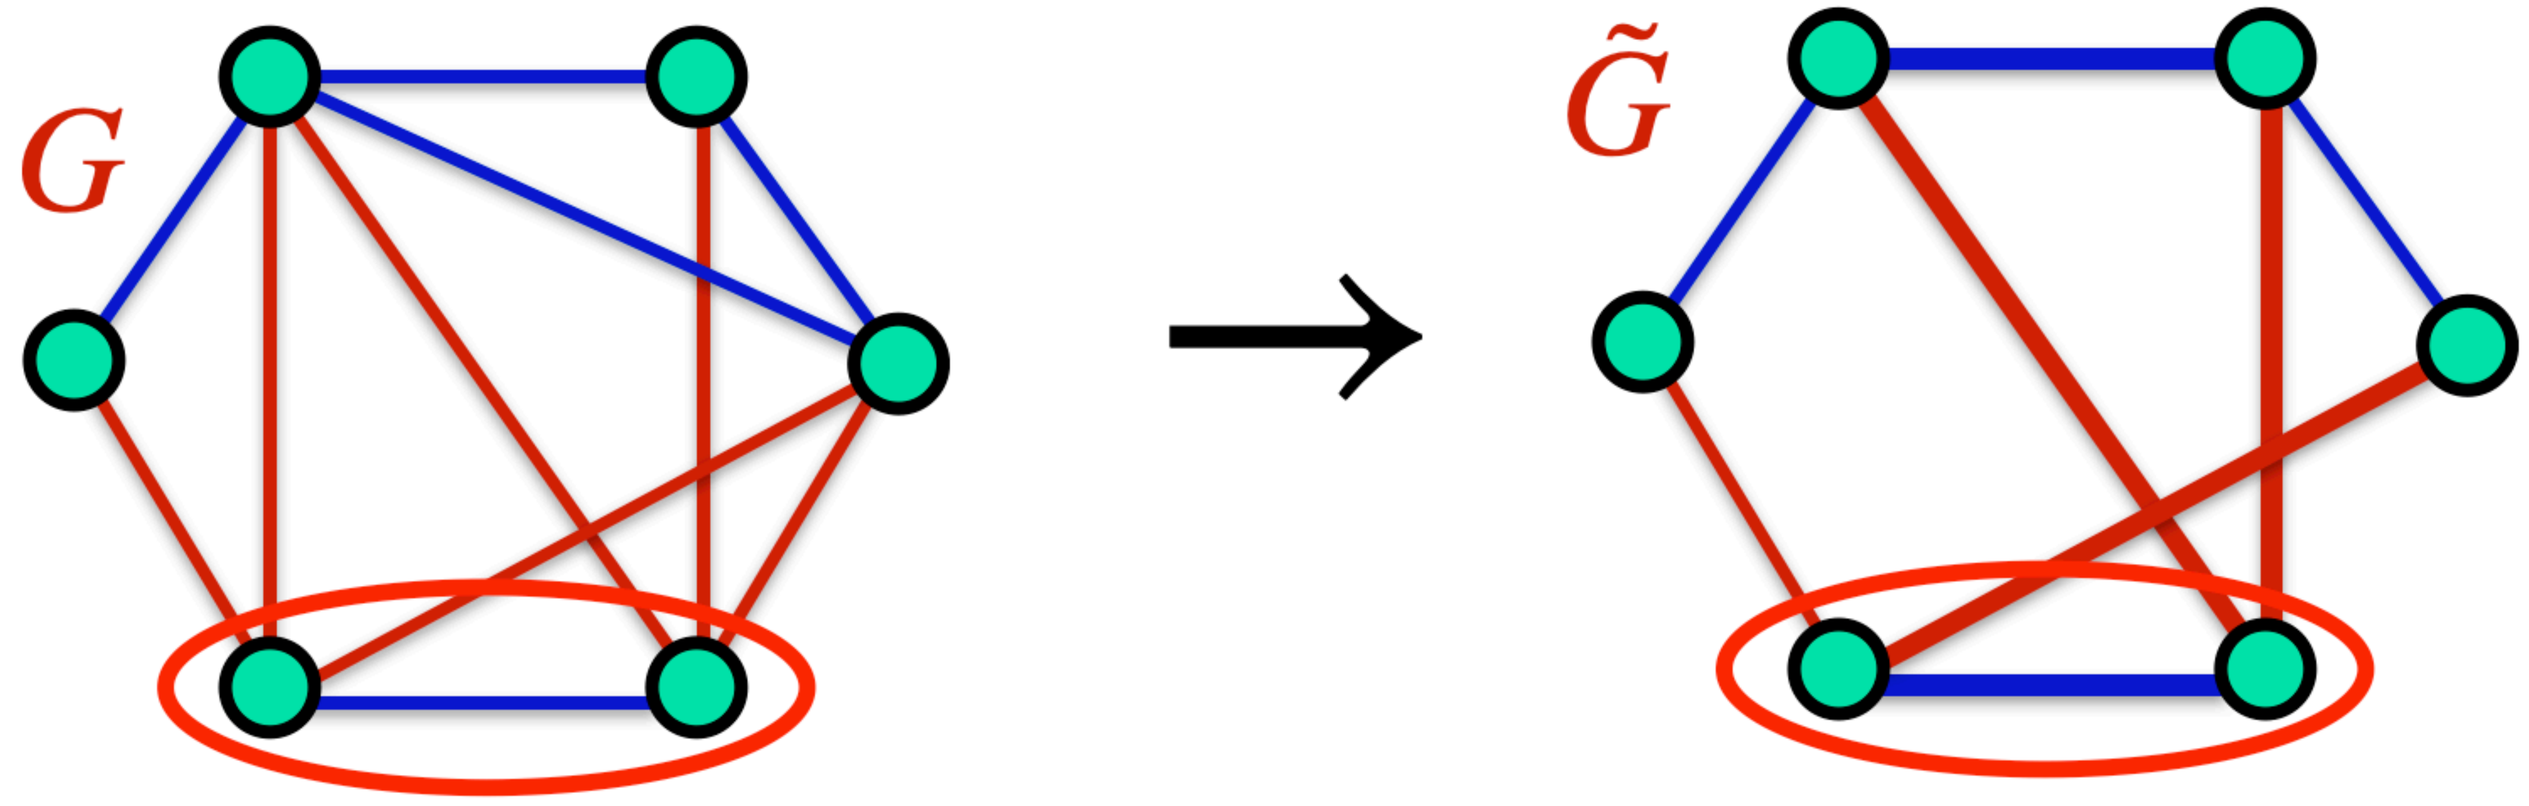

Figure due to Aaron Putterman

Spectral Sparsifiers

- Cut sparsifiers have a **spectral** generalization
- For any graph $G = G(V, E)$, the **Laplacian** L_G defines the quadratic form

$$\langle x, L_G x \rangle := \sum_{e=\{u,v\} \in E} w_e (x_u - x_v)^2$$

- Here $x \in \mathbb{R}^n$, w_e is weight of edge e
- For $x = \mathbf{1}_S$, $\langle x, L_G x \rangle = \mathrm{cut}_G(S)$

Spectral Sparsifiers

- For a graph $G = G(V, E)$, a spectral sparsifier is a sparse graph $\tilde{G} = \tilde{G}(V, \tilde{E})$ such that for every $x \in \mathbb{R}^n$ we have

$$\langle x, L_{\tilde{G}} x \rangle \in [0.99, 1.01] \cdot \langle x, L_G x \rangle$$

- **Putting $x = \mathbf{1}_S$ recovers cut sparsification! More general problem.**

- Notation: If $\langle x, Ax \rangle \leqslant \langle x, Bx \rangle$ for all x, we say $A \preccurlyeq B$ (a.k.a Loewner order)

- Thus **spectral sparsification** $\iff 0.99 L_G \preccurlyeq L_{\tilde{G}} \preccurlyeq 1.01 L_G$

Existence of Spectral Sparsifiers

- Spectral sparsifiers of size $O(n)$ exist, and can be constructed in $\text{poly}(n)$ time: Works by Benczur-Karger, Spielman-Teng, Spielman-Srivastava, Batson-Spielman-Srivastava

- Lots of applications in Laplacian solvers, graph algorithms, etc.

- $O(n)$ is the optimal spectral sparsifier size!
  - Sparsifiers of connected graphs must be connected
  - Need $\geq n-1$ edges to be connected!

Plan for the Talk

- **Quantum Sparsification:** Will discuss the quantum generalization of graph sparsification, quantum cut sparsifiers
  - Will obtain “quantum cut” sparsifiers of size $\tilde{O}(n)$
  - Quantum generalization of Spielman-Teng, Spielman-Srivastava!

Quantum-Cut Sparsifiers

- What is quantum cut?
  - What is *classical* cut?

- For any graph $G = G(V, E)$ and vector $x \in \mathbb{R}^n$, the classical cut corresponding to edge $e = \{u, v\}$ is
$$\langle x, L_e x \rangle := (x_u - x_v)^2.$$
  - Suppose $x = \mathbf{1}_S$. Suppose $u \in S, v \notin S$. Then $(x_u - x_v)^2 = (1 - 0)^2 = 1$.
  - Note that $L_G = \sum_e w_e L_e$
  - Classical cuts "act" on vectors $x \in \mathbb{R}^n \cong \mathbb{R}^V$
  - Quantum cuts "act" on vectors $\psi \in \mathbb{R}^{2^n} \cong \mathbb{R}^{2^V}$. Recall $2^V = \{S : S \subseteq V\}$.

Quantum Cut: Definition

- For any $\psi \in \mathbb{R}^{2^V}$, and edge $e = \{u, v\}$

$$\langle \psi, \mathcal{L}_e \psi \rangle := \frac{1}{2} \cdot \sum_{S \in 2^V : |S \cap \{u,v\}| = 1} (\psi(S) - \psi(S \oplus \{u, v\}))^2$$

- We only care about S which contain **exactly** one of u or v, i.e. $|S \cap \{u, v\}| = 1$
- For those S, $S \oplus \{u, v\}$ **swaps** u and v in S. Note that $|S| = |S \oplus \{u, v\}|$
- The $\frac{1}{2}$ is because $(S, S \oplus e)$ also appears in the sum as $(S \oplus e, (S \oplus e) \oplus e) = (S \oplus e, S)$

Quantum Cut: Example

- Example: Suppose $V = \{1,2,3\}$. Suppose $e = \{1,2\}$ with $u = 1, v = 2$
- Only sets S which matter are $S = \{1\}, \{2\}, \{1,3\}, \{2,3\}$
- For these S, $S \oplus \{u, v\}$ **swaps out** u, v
- $\{1\} \leftrightarrow \{2\}, \{1,3\} \leftrightarrow \{2,3\}$

Quantum Cut: Example

$$\langle \psi, \mathscr{L}_e \psi \rangle = \frac{1}{2} \cdot \sum_{S \in 2^V : |S \cap \{u,v\}| = 1} (\psi(S) - \psi(S \oplus \{u, v\}))^2$$

- **Example**: Suppose $V = \{1,2,3\}, e = \{1,2\}$. ψ is a $2^3 = 8$-dimensional vector.
  - $\{1\} \leftrightarrow \{2\}, \{1,3\} \leftrightarrow \{2,3\}$
  - Thus

$$\langle \psi, \mathscr{L}_e \psi \rangle = (\psi(\{1\}) - \psi(\{2\}))^2 + (\psi(\{1,3\}) - \psi(\{2,3\}))^2$$
$$= (1 - 0)^2 + (2 - 1)^2 = 2$$

{}	-2
{1}	1
{2}	0
{3}	-1
{1, 2}	3
{1, 3}	2
{2, 3}	1
{1, 2, 3}	5

For those in the know…

- Some of you might have seen the definition of quantum cut as

$$(|01\rangle - |10\rangle)(\langle 01| - \langle 10|)_e \otimes \mathrm{Id}_{[n]\backslash e}$$

- This is equivalent to the definition that I just showed you!

- We will continue to work with our definition as it is more convenient!

Quantum Laplacians

$$\langle \psi, \mathscr{L}_e \psi \rangle = \frac{1}{2} \cdot \sum_{S \in 2^V : |S \cap \{u,v\}| = 1} (\psi(S) - \psi(S \oplus \{u, v\}))^2$$

- As in the classical case, for a (weighted) graph $G = G(V, E)$ define the **quantum Laplacian**

$$\mathscr{L}_G := \sum_e w_e \mathscr{L}_e$$

- $\mathscr{L}_G$ is a $2^n \times 2^n$ matrix! Recall the **classical Laplacian** L_G was a $n \times n$ matrix

Quantum Cut Sparsification

- For a graph $G = G(V, E)$, a spectral sparsifier is a sparse graph $\tilde{G} = \tilde{G}(V, \tilde{E})$ such that for every $x \in \mathbb{R}^n$ we have
$$\langle x, L_{\tilde{G}} x \rangle \in [0.99, 1.01] \cdot \langle x, L_G x \rangle \iff 0.99 L_G \preccurlyeq L_{\tilde{G}} \preccurlyeq 1.01 L_G$$

- For a graph $G = G(V, E)$, a quantum-cut sparsifier is a sparse graph $\tilde{G} = \tilde{G}(V, \tilde{E})$ such that for every $\psi \in \mathbb{R}^{2^n}$ we have
$$\langle \psi, \mathscr{L}_{\tilde{G}} \psi \rangle \in [0.99, 1.01] \cdot \langle \psi, \mathscr{L}_G \psi \rangle \iff 0.99 \mathscr{L}_G \preccurlyeq \mathscr{L}_{\tilde{G}} \preccurlyeq 1.01 \mathscr{L}_G$$

Connections

- Quantum cut studied in the quantum computing community as analog of classical cut
  - Like max-cut, (approx.) algos for **quantum** max cut well-studied
  - Also studied in physics as **anti-ferromagnetic Heisenberg model**

- The process $S \leftrightarrow S \oplus \{u, v\}$ for S such that $|S \cap \{u, v\}| = 1$ is a.k.a. **symmetric exclusion process** in probability theory. More on this later

- For any $e = \{u, v\}$, consider the graph on 2^V where

$$S \text{ is connected to } S \oplus \{u, v\} \text{ if } |S \cap \{u, v\}| = 1$$

- This graph is known as the **Kikuchi graph** $\mathscr{K}_e$ generated by e. Also write $\mathscr{K}_G := \bigcup_e \mathscr{K}_e$
  - Kikuchi graphs show up in tensor PCA, extremal combinatorics, random CSPs, …

Quantum Sparsification harder than Classical

- Recall that for $\psi \in \mathbb{R}^{2^V}$,

$$\langle \psi, \mathscr{L}_e \psi \rangle = \frac{1}{2} \cdot \sum_{S : |S \cap e| = 1} (\psi(S) - \psi(S \oplus e))^2 .$$

- Suppose $\psi(S) = 0$ for all S with $|S| \neq 1$.
- Then the only sets S which contribute to $\langle \psi, \mathscr{L}_e \psi \rangle$ are $\{u\}, \{v\}$
- Thus $\langle \psi, \mathscr{L}_e \psi \rangle = (\psi(\{u\}) - \psi(\{v\}))^2$
- But this is exactly "classical" cut!
- Quantum sparsification requires us to preserve $\langle \psi, \mathscr{L}_G \psi \rangle$ for **all** ψ
- But for ψ supported on singleton sets, quantum and classical cuts coincide!
- Thus quantum sparsification more general problem than "classical" spectral sparsification

A first attempt: The Chernoff Problem

- Thus quantum sparsifiers need $\Omega(n)$ edges, since "classical" spectral sparsifiers also need that many

- $O(n)$ edges suffice for spectral sparsifiers though
  - We will see a simpler proof of $\tilde{O}(n)$ sized sparsifiers

- Let $\{L_e\}_{e \in E(G)}$ be the edge Laplacians of our graph, and let $L_G = \sum_e L_e$ be the graph Laplacian
  - Let $\tilde{L}_e = L_G^{-1/2} L_e L_G^{-1/2}$ be the "reduced" edge Laplacian, such that $\sum_e \tilde{L}_e = \mathrm{Id}$
  - Yes, the inverse should really be a pseudo-inverse, but we ignore the difference

A first attempt: The Chernoff Problem

- **Matrix Concentration** tells us that we can get sparsifiers of size $\log \dim(L_e) \cdot \sum_e \|\tilde{L}_e\|_2$
- Can show that $\sum_e \|\tilde{L}_e\|_2 \leq n$. Since $\log \dim(L_e) = \log(n)$, we get sparsifier of size $O(n \log(n))$, as desired.
- What about quantum Laplacians $\mathscr{L}_e$ and their reduced counterparts $\tilde{\mathscr{L}}_e := \mathscr{L}_G^{-1/2} \mathscr{L}_e \mathscr{L}_G^{-1/2}$?

A first attempt: The Chernoff Problem

- What about quantum Laplacians $\mathscr{L}_e$ and their reduced counterparts $\tilde{\mathscr{L}}_e := \mathscr{L}_G^{-1/2} \mathscr{L}_e \mathscr{L}_G^{-1/2}$?
- Can still show that $\sum_e \|\tilde{\mathscr{L}}_e\|_2 \leq \tilde{O}(n)$!
  - Proof non-trivial, but not difficult
- So we obtain sparsifiers of size $\log \dim(\tilde{\mathscr{L}}_e) \cdot \sum_e \|\tilde{\mathscr{L}}_e\|_2$
- But $\log \dim(\tilde{\mathscr{L}}_e) = \log(2^n) = n$. Thus we obtain sparsifier of size $n \cdot \tilde{O}(n) = \tilde{O}(n^2)$

A first attempt: The Chernoff Problem

- $\tilde{O}(n^2)$ is a worse than useless bound, since a graph can have at most $O(n^2)$ edges anyways!

- Need new ideas! $\log(2^n)$ too expensive!

Main Idea 1: Degree Decomposition

- We still want to apply the Matrix Chernoff bound, but losing $\log(2^n)$ is very expensive

- Can we chop up our space into “smaller” pieces, and not lose as much in the $\log(\text{dimension})$ terms?

- Yes: The Degree Decomposition!

Key Idea 1: Degree Decomposition

- There exists a decomposition

$$\mathbb{R}^{2^n} = \bigoplus_{j=0}^{n/2} V_j$$

of $\mathbb{R}^{2^n}$ such that for any $e = \{u, v\}$ we have $\mathscr{L}_e(V_j) \subseteq V_j$, and

$$\dim(V_j) \leq n^{j+2}\,.$$

- In other words, there is a **graded decomposition** which every quantum Laplacian $\mathscr{L}_e$ respects

- Where do these V_js come from? V_j is the subspace of all "degree-j" functions $2^V \to \mathbb{R}$

Degree of a function

- There is a natural partition of $2^{[n]} \cong \{0,1\}^n$, the slices $W_\ell := \{S \in 2^{[n]} : |S| = \ell\}$

- A function $W_0 \to \mathbb{R}$ is a "degree 0" function
  - Note that a function $W_0 \to \mathbb{R}$ is a constant function, because W_0 has size 1!
  - We thus declare all constant functions $W_\ell \to \mathbb{R}$ to be degree 0.

- Now take a function $f : W_1 \to \mathbb{R}$

Higher Degree Functions

- We declare all constant functions $W_\ell \to \mathbb{R}$ to be degree 0.
- Now take a function $f : W_1 \to \mathbb{R}$
  - Remove the "constant" component of f, i.e. consider $g = f - \mathbb{E}[f]$. Note that $\mathbb{E}[g] = 0$
  - After removing the degree 0 component from a function on W_1, we are left with a degree 1 function!

- We can keep going!

- Take a function $f : W_2 \to \mathbb{R}$

Higher Degree Functions

- Take a function $f : W_2 \to \mathbb{R}$
- A degree 0 function on W_2 is simply a constant function
- What does a "degree 1" function on W_2 look like? It looks like $g(\{i, j\}) = h(i) + h(j)$, i.e. it looks at coordinates "one at a time", and is symmetric b/w i, j
- Consider the orthogonal complement of degree 0 and degree 1 functions on W_2
- What remain are the degree 2 functions!

Higher Degree Functions

- A degree 0 function is simply a constant function
- A "degree 1" function looks like $g(\{i,j\}) = h(i) + h(j)$
- A degree 2 function looks like $g(\{i,j,k\}) = h(\{i,j\}) + h(\{j,k\}) + h(\{k,i\})$
- If you take orthogonal complement of these functions $W_3 \to \mathbb{R}$, what remain are the degree 3 functions!

What is a degree j function?

- A function $\psi : 2^{[n]} \to \mathbb{R}$ can be viewed as a collection of functions $\psi_\ell : W_\ell \to \mathbb{R}$
  - W_ℓ is isomorphic to $W_{n-\ell}$ via $S \leftrightarrow [n]\backslash S$, and thus we only need to consider $0 \leq \ell \leq n/2$

- If each ψ_ℓ is either 0 or degree j we call ψ a degree j function

- It is not hard to see that the subspace of degree j functions have dimension $\approx n^j$

Why do quantum Laplacians preserve degree?

- Why does $\mathscr{L}_e$ preserve degree?

- $\mathscr{L}_e$ swaps/relabels coordinates, hence can not change degree!

- **Claim:** $\mathscr{L}_e = \mathrm{Id} - \tau_e$ where $(\mathrm{Id} \cdot \psi) := \psi$, and

$$(\tau_e \cdot \psi)(x_1, \ldots, x_u, \ldots, x_v, \ldots, x_n) := \psi(x_1, \ldots, x_v, \ldots, x_u, \ldots, x_n)$$

where $S \in 2^V$ can be written as a $0-1$ indicator vector $\mathbf{1}_S = (x_1, \ldots, x_n)$.

- $\tau_e = \tau_{u,v}$ simply swaps the u^{th} and v^{th} coordinates

- Swapping/relabeling coordinates can not change the degree of a function!

Quantum Laplacians and Swaps

- **Claim:** $\mathscr{L}_e = \mathrm{Id} - \tau_e$ where $(\mathrm{Id} \cdot \psi) := \psi$, and

$$(\tau_e \cdot \psi)(x_1, \ldots, x_u, \ldots, x_v, \ldots, x_n) := \psi(x_1, \ldots, x_v, \ldots, x_u, \ldots, x_n)$$

- $\tau_e = \tau_{u,v}$ simply swaps the uth and vth coordinates

- Thus $(\mathscr{L}_e \psi)(S) = \psi(S) - \psi(S$ with u, v swapped).

- If S does not contain u or v, then (S with u, v swapped) $= S$
- If S contains both u and v, then (S with u, v swapped) $= S$
- If S contains exactly one of u or v, then (S with u, v swapped) $= S \oplus \{u, v\}$

- Hence

$$\langle \psi, \mathscr{L}_e \psi \rangle = 0.5 \sum_{S : |S \cap e| = 1} (\psi(S) - \psi(S \oplus e))^2$$

Quantum Laplacians and Swaps

- Hence

$$\langle \psi, \mathscr{L}_e \psi \rangle = 0.5 \sum_{S:|S \cap e|=1} (\psi(S) - \psi(S \oplus e))^2$$

- Thus $\mathscr{L}_e = \mathrm{Id} - \tau_e$, and thus $\mathscr{L}_e$ preserves degree since neither Id nor τ_e change degree

What have we learnt so far?

- There exists a decomposition

$$\mathbb{R}^{2^n} = \bigoplus_{j=0}^{n/2} V_j$$

of $\mathbb{R}^{2^n}$ such that for any $e = \{u, v\}$ we have $\mathscr{L}_e(V_j) \subseteq V_j$, and

$$\dim(V_j) \leq n^{j+2}.$$

- Since all Laplacians respect this decomposition, we can restrict our attention to V_j, and in the end we can union bound over $n/2$ different $V_j s$
  - Union bounding over $n/2$ things can be done by losing a $\log(n)$ factor

Will Random Sampling Work?

- Thus let's focus on V_j, which has dimension $n^{O(j)}$ (assuming $j \leq n/2$)
- Suppose we want to sparsify $\mathscr{L}_G$. Focus on $\mathscr{L}_G|_{V_j}$
- But even in this space, Matrix Chernoff loses $\log \dim(V_j) = \log(n^{O(j)}) = \tilde{O}(j)$, which is bad for $j \geq \Omega(n)$. Have to do something cleverer.

Will Random Sampling Work?

- Suppose $\tilde{G}$ is a **random sample** of G, i.e. every edge in G is retained in $\tilde{G}$ with equal probability, and suppose $\tilde{G}$ has $n\,\mathrm{polylog}(n)$ many edges w.h.p
  - Recall that we want a sparsifier of size $n\,\mathrm{polylog}(n)$, so if $\tilde{G}$ sparsifies G we are very happy

- **Matrix Concentration: $\tilde{G}$ fails to sparsify G** with probability

$$\lesssim \dim(V_j) \cdot n^{-\lambda_{\min}(\mathscr{L}_G|_{V_j})/d_{\text{avg}}(G)}$$

where $d_{\text{avg}}(G)$ is the average degree of G

- Thus can kill $\dim(V_j)$ if $\lambda_{\min}(\mathscr{L}_G|_{V_j}) \geq \Omega(j \cdot d_{\text{avg}}(G))$

Why does $\lambda_{\min}(\mathscr{L}_G|_{V_j})$ need to grow as $j \cdot d_{\text{avg}}(G)$?

- Intuitively, if $\mathscr{L}_G|_{V_j}$ has large eigenvalues, obtaining concentration is easier. Concentration hardest to obtain when $\mathscr{L}_G|_{V_j}$ has small eigenvalues
  - This is because we want **multiplicative** guarantees of the form $0.99\mathscr{L}_G \preceq \mathscr{L}_{\tilde{G}} \preceq 1.01\mathscr{L}_G$
  - Thus if $\mathscr{L}_G$ has small eigenvalues, room for error also smaller
  - Equivalently, larger eigenvalues allow us to apply matrix conc. over larger spaces
- But do we have $\lambda_{\min}(\mathscr{L}_G|_{V_j})$ growing as $j \cdot d_{\text{avg}}(G)$ for all graphs G? For any graph G?

For Complete Graphs, Yes!

- **Need:** If $\lambda_{\min}(\mathcal{L}_G|_{V_j}) \gg j \cdot d_{\text{avg}}(G)$, then $0.99\mathcal{L}_G \preceq \mathcal{L}_{\tilde{G}} \preceq 1.01\mathcal{L}_G$
  - Is this achievable?
- For G being the complete graph, yes!
- When G is the complete graph, can exactly compute all eigenvalues of $\mathcal{L}_G|_{V_j}$
- All eigenvalues are equal to $j(n+1-j)$, which is $\geq \Omega(jn) = \Omega(j \cdot d_{\text{avg}}(G))$

What about other graphs?

- For G being the complete graph, $\lambda_{\min}(\mathscr{L}_G|_{V_j}) \geq \Omega(jn)$
  - What about other graphs? What about expanders?

- Unfortunately, direct computation fails, because no other graph, (except the empty one) is as symmetric as the complete one

- Enter the second main idea: The **Alon-Kozma inequality**

Main Idea 2: The Alon-Kozma Inequality

- Let G be a d-regular expander with $1/\mathrm{polylog}(n)$ spectral gap
  - Such expanders can be obtained very generically through expander decompositions
- Then $L_G \succeq \tilde{\Omega}(d/n) \cdot L_{K_n}$, where K_n is the complete graph. This is simply the definition of spectral gap. L_G is the usual "classical" graph Laplacian
- **Alon-Kozma inequality**: The same inequality transfers over to the quantum Laplacians!

$$\mathscr{L}_G \succeq \tilde{\Omega}(d/n) \cdot \mathscr{L}_{K_n}$$

- A comparison between two $n \times n$ operators gave us a comparison between two $2^n \times 2^n$ operators!

Quantum-Cut Sparsifiers using Alon-Kozma

- First suppose G is an expander. Then by Alon-Kozma, we have

$$\lambda_{\min}(\mathscr{L}_G|_{V_j}) \geq \tilde{\Omega}(1) \cdot (d_{\text{avg}}(G)/n) \cdot \lambda_{\min}(\mathscr{L}_{K_n}|_{V_j})$$
$$= \tilde{\Omega}(1) \cdot (d_{\text{avg}}(G)/n) \cdot jn = \tilde{\Omega}(j \cdot d_{\text{avg}}(G))$$

- Thus **random sampling sparsifies** G

- For general graphs G, run an expander decomposition to write $G =$ Expander $+$ Residual

Expander Decompositions

- Thus **random sampling sparsifies expander graphs** G

- For general graphs G, run an **expander decomposition** to write $G = \text{Expander} \ + \ \text{Residual}$
  - For every graph G with $\tilde{\Omega}(n)$ many edges, can write G as a union of expanders + a "Residual" graph, such that the number of edges in "Residual" graph is $\leq 1\%$ of edges in G
  - Can be done very efficiently, see, for e.g., Saranurak-Wang'19

Quantum Cut Sparsification

- Thus **random sampling sparsifies expander graphs** G

- For arbitrary graphs G, write $G =$ Expander $+$ Residual

- Sparsify the “Expander” down to $\tilde{O}(n)$ many edges using random sampling

- Apply expander decomposition again on the Residual graph!

- Only $\tilde{O}(1)$ many rounds to this process, since #edges in Residual graph is $\leq 1\,\%$ of #edges in G

Summary of our Ideas

- Run **degree decomposition** to obtain spaces V_j such that all quantum Laplacians "respect" V_j, i.e. $\mathscr{L}_e(V_j) \subseteq V_j$
- $\dim(V_j)$ grows with j, so need to show that $\lambda_{\min}(\mathscr{L}_G|_{V_j})$ grows equally fast so that matrix Chernoff can be applied
- Verify that the growth condition on $\lambda_{\min}(\mathscr{L}_G|_{V_j})$ holds for G being the complete graph
- Transfer the growth condition from complete graphs to expanders using Alon-Kozma
- Iteratively run (expander decomposition + sparsification) on general graphs
- **QED!**

The Magic Ingredient: Alon-Kozma Inequality

- The Alon-Kozma inequality is the key ingredient which makes this proof work

- It was established by Alon and Kozma in the context of *interchange processes* in probability theory

- The interchange process is the following:
  - Suppose we have n distinguishable particles on the vertices of our graph
  - At every instance, we pick a random edge, and swap the particles on its endpoints
  - “Configuration” of the system described by a permutation on n particles

- Alon-Kozma showed that the spectral gap of the graph translates to the spectral gap of this Markov chain

The Magic Ingredient: Alon-Kozma Inequality

- Alon-Kozma showed that the spectral gap of the graph translates to the spectral gap of this Markov chain
- Deriving what we need from this is easy
- We recall that when quantum Laplacians send S to $S \oplus \{u, v\}$, we have $|S| = |S \oplus \{u, v\}|$ since $|S \cap \{u, v\}| = 1$
- Thus we can imagine S as consisting of $k = |S|$ particles, and our quantum Laplacian **swaps** a particle along $\{u, v\}$
- This process is a.k.a. the k-**particle exclusion process**
- It is a sub-process of the **interchange process**: Why? Imagine instead of all n particles being distinguishable, we had k red particles, and $n - k$ blue particles.
  - It would behave like the exclusion process
  - Thus the Alon-Kozma spectral gap transfer statement holds

Thanks!

Find our paper here!
arxiv.org/abs/2606.09728