A Simple Framework for Finding Balanced Sparse Cuts via APSP

Li Chen (Georgia Tech)

SOSA 2023

Joint work with

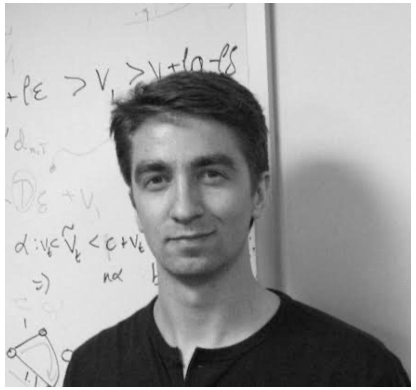

Rasmus Kyng
ETH

Maximilian
Probst Gutenberg
ETH

Sushant
Sachdeva
U. Toronto

Sparse Cuts and Expanders

- Graph $G = (V, E)$, a cut $S \subseteq V$,
- **Sparsity** of a cut:

$$\Psi(S) = \frac{E(S,\ V \setminus S)}{\min\{|S|,\ |V \setminus S|\}}$$

- A cut is **balanced** if $\min\{|S|,\ |V \setminus S|\} = \Omega(n)$
- G is a $\boldsymbol{\psi}$**-expander** if every cut has sparsity at least ψ

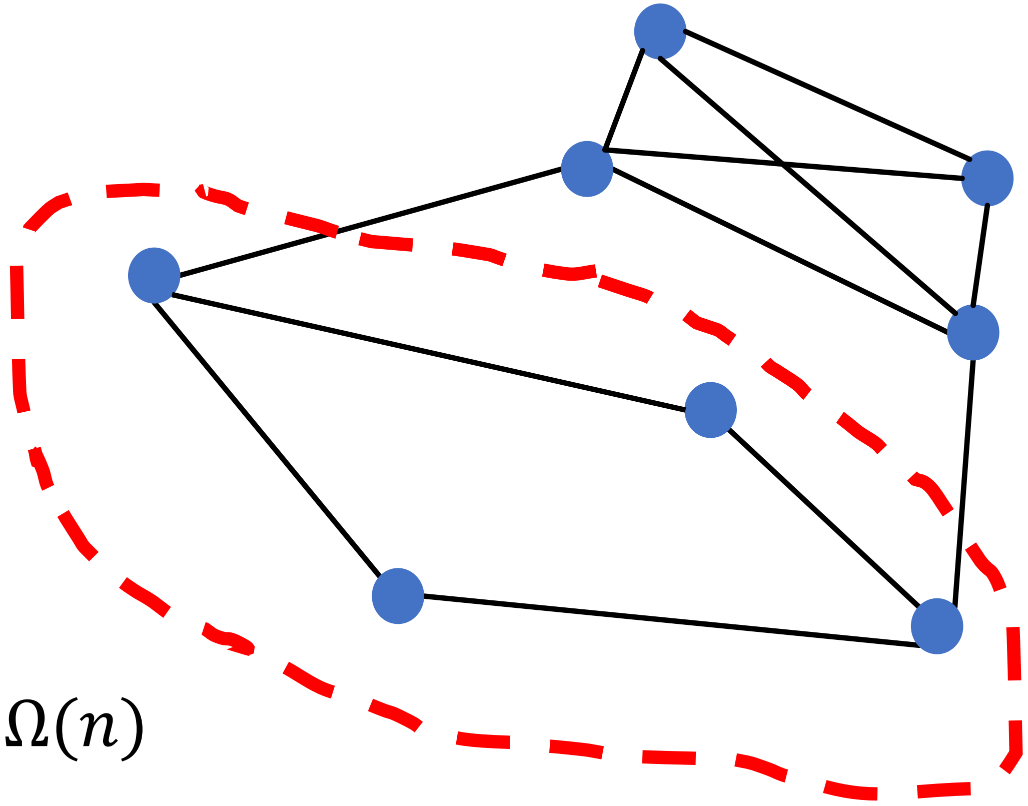

Sparsity = 2 / 4 = 0.5

Balanced Sparse Cut Problem

- Given graph $G = (V, E)$ of max deg 10

Assume WLOG

- Sparsity ψ, approx ratio $\alpha \geq 1$
- ***Find*** a cut $S \subseteq V$ with $\min\{|S|,\ |V \setminus S|\} \geq n/100$ of sparsity $\leq \psi$
- Or, ***certify*** that every cut $\mathrm{X} \subseteq V$ with $\min\{|X|,\ |V \setminus X|\} = \Omega(n)$ has sparsity $\geq \psi/\alpha$

APX-hard, i.e. $\alpha > 1$ for P

Assuming UGC, $\alpha = \Omega(1)$ for P

Applications

- Graph Clustering
- Expander Decomposition
  - Max Flow
  - All-Pair Min-Cut
  - Laplacian Systems
  - Dynamic Connectivity
  - …

Our Result

- Given graph $G = (V, E)$ of max deg 10, sparsity ψ
- ***Find*** a cut $S \subseteq V$ with $\min\{|S|,\ |V \setminus S|\} \geq n/100$ of sparsity $\leq \psi$
- Or, ***certify*** that every cut $\mathrm{X} \subseteq V$ with $\min\{|X|,\ |V \setminus X|\} = \Omega(n)$ has sparsity $\psi/O(\log n \log\log n)$
- The algorithm runs in $\tilde{O}(n^2/\psi)$ time, in particular, $O(n)$ shortest path computations

Our Result

- Given graph $G = (V, E)$ of max deg 10, sparsity ψ
- ***Find*** a cut $S \subseteq V$ with $\min\{|S|,\ |V \setminus S|\} \geq n/100$ of sparsity $\leq \psi$
- Or, ***certify*** that every cut $X \subseteq V$ with $\min\{|X|,\ |V \setminus X|\} = \Omega(n)$ has sparsity $\psi/\boldsymbol{O(\log^2 n)}$
- The algorithm runs in $\tilde{O}(n^2/\psi)$ time, in particular, $O(n)$ shortest path computations

Embeddings

- graphs $G = (V, E), H = (V, F)$
- **Embedding** $\Pi_{H\to G}$ is a collection of path in G
  - $\forall\ (u, v) \in H, \Pi_{H\to G}$ contains a uv-path in G
- **Congestion** $cong(\Pi_{H\to G}) = c$ if every edge in G is used $\leq c$ times

Certificate: Almost Embed an Expander

Let H be a ψ-expander and $H' \subseteq H$ has all but $\psi n/200$ edges.

If $cong(\Pi_{H' \to G}) = c$, G has no balanced $O(\psi/c)$-sparse cuts.

- Let S be any balanced cut, i.e., $n/100 \leq |S| \leq n/2$
- $|E_{H'}(S, V \backslash S)| \geq |E_H(S, V \backslash S)| - E(H \backslash H') \geq \psi|S| - \psi n/200 \geq \psi|S|/2$
- Every cut edge in H' uses a path crossing the cut in G at least once.
- Every cut edge in G is used at most c times.
- $|E_G(S, V \backslash S)| \geq |E_{H'}(S, V \backslash S)|/c \geq \psi|S|/2c$

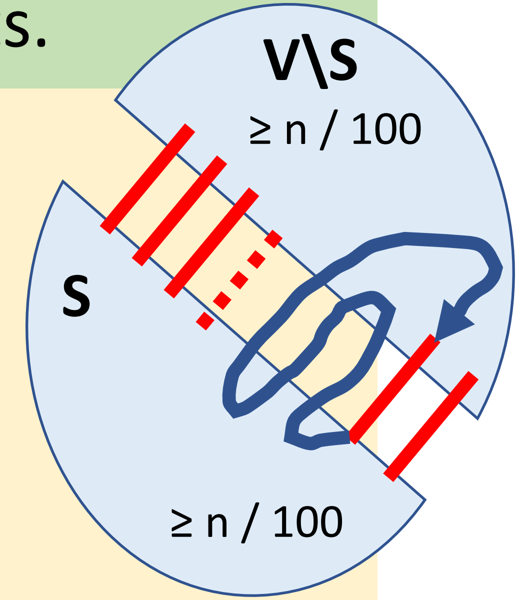

Plan

- Embed an $\Omega(1)$-expander with **short** paths and **low** congestion
- If can embed most edges, G has **no** balanced sparse cut
- Otherwise, try to find a balanced sparse cut
- Multiplicative Weight Update (MWU) to control congestion

Alg: Embed Expander via Shortest Paths

- Initialize edge weights $w_e = 1$ on G. Step size $\eta = \psi / \log n$.
- Let H be a $\Omega(1)$-expander. $H' = \emptyset$
- For each edge $e = (u, v) \in H$ s.t. $dist_{G,w}(u, v) = O(\log n / \psi)$
  - Add e to H'
  - Embed e via P, the shortest uv-path in G
  - For each $f \in P$, update $w_e = w_e(1 + \eta)$

Penalize congested edges

- **FAIL**: If $H \backslash H'$ has more than $n/100$ edges
- **SUCCESS**: Output the embedding $\Pi_{H' \to G}$ and H'

Success: embed all but few expander edges

- H' has all but $n/100$ edges.
- Every balanced cut of G has sparsity at least $\Omega(1/cong)$.

Success: embed all but few expander edges

- H' has all but $n/100$ edges.
- Every balanced cut of G has sparsity at least $\Omega(\boldsymbol{\psi}/\,\boldsymbol{log^2 n})$.

$$cong = O(\log^2 n/\psi)$$

- Embed $\mathrm{e} = (u, v)$ with path P, $\Delta\|w\|_1 = \eta \cdot w(P) = O(\eta \log n\,/\psi)$
- Each edge in H' increase $\|w\|_1$ by $O(1)$, final $\|w\|_1 = O(n)$
- Edge f is used c times, $w_f = (1 + \eta)^c = O(n)$

- For each $f \in P$, update $w_e = w_e(1 + \eta)$

- $c = O(\log n\,/\eta) = O(\log^2 n/\psi)$

Fail: many distant pairs

- $F = H\backslash H'$ contains $n/100$ pairs each of which are $\Omega(\log n / \psi)$ apart.

If u and v are D apart, can find cut of sparsity $O(\log n / D)$

- Ball-Growing! Define $B_i = \{x | dist(u, x) \le i\}, i = 0,1, \dots, D/2$
- $\exists i$: $vol(B_{i+1}) < \left(1 + \frac{100 \log n}{D}\right) vol(B_i)$
- $|E(B_i, V\backslash B_i)| \le vol(B_{i+1}) - vol(B_i)$

$$< \frac{100 \log n}{D} vol(B_i) < \frac{1000 \log n}{D} |B_i|$$

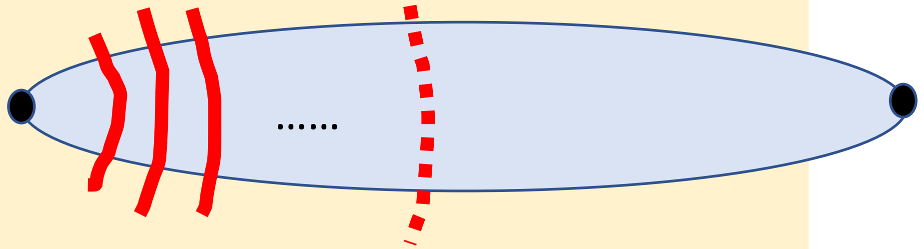

- Similar argument holds in our case (weighted with total weight $\|w\|_1 = O(n)$)

Fail: many distant pairs

- $F = H \backslash H'$ contains $n/100$ pairs each of which are $\Omega(\log n \,/\psi)$ apart.
- For each pair, find a ψ-sparse cut and peel the smaller side
- Final cut is ψ-sparse and has at least $n/100$ vertices on each side

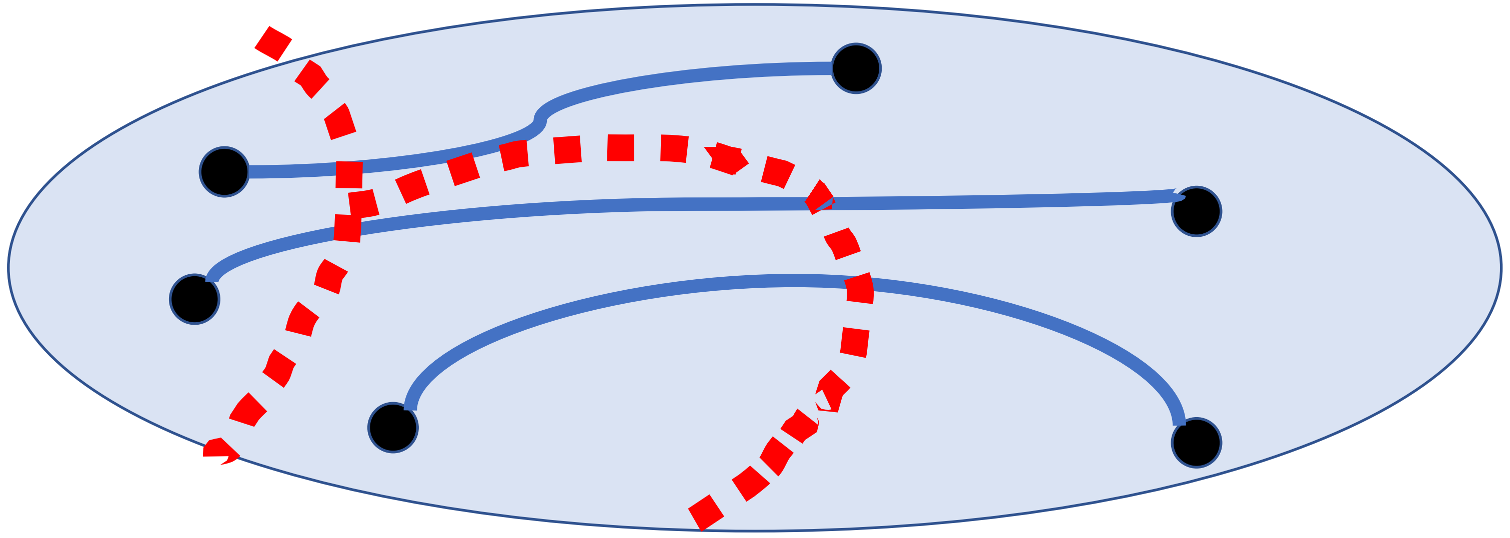

Conclusion

- $O(\log n \log\log n)$-Balanced Sparse Cut via $O(n)$ Shortest Paths
  - Sparsity or Conductance (arbitrary degree)
- Compute Shortest Paths on graph w/ increasing weight
- κ-approx SP gives $O(\kappa \log n \log\log n)$-approx. balanced sparse cut
  - $m^{1+o(1)}$-time/ $m^{o(1)}$-approx Decremental APSP gives $m^{1+o(1)}/\psi$ runtime
    [Chuzhoy '21, Bernstein-Probst Gutenberg-Saranurak '21]
- Approx. multi-comm flow from MWU to decremental APSP

Thanks!!

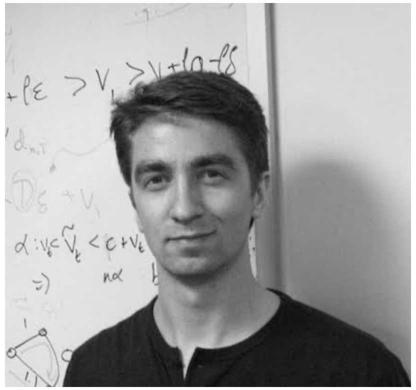

Rasmus Kyng
ETH

Maximilian
Probst Gutenberg
ETH

Sushant
Sachdeva
U. Toronto