Low-Congestion Shortcuts

9/6/2021
Lecture 9

Scribe: Avi Cohen

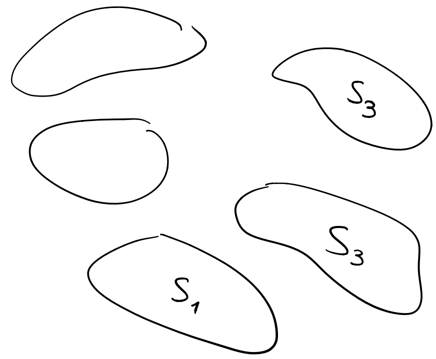

Part-wise Aggregation

Collection of vertex disjoint $S_1, \dots, S_\ell$

each $G[S_i]$ is connected.

Every vertex $v_i \in S_j$ holds $O(\log n)$-bit x_i

and it is reqired for every $v_i \in S_i$ to learn

the aggregate function of all x_k's for $v_k \in S_j$.

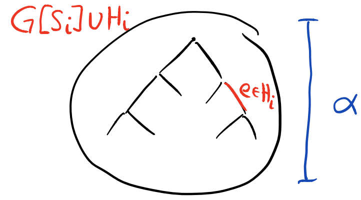

Low-Congestion Shortcuts (LCS)

$G=(V,E)$ and $S_1, \dots, S_\ell$, each $G[S_i]$ is connected.

A collection of subgraphs $\mathcal{H} = \{H_1, \dots, H_\ell\}$ is an (α, β) LCS if:

1. $\mathrm{Diam}(G[S_i] \cup H_i) \le \alpha$.
2. Each edge $e \in G$ appears on at most β H_i-subgraphs.

<u>Theorem</u>: Given a T-round alg for computing (α, β) shortcuts for $S_1, \dots, S_\ell$, then one can solve the partwise aggregation problem in $\tilde{O}(T + \underbrace{\alpha + \beta}_{\text{The Quality of the Shortcut.}})$.

<u>Remark</u>: This also implies MST in $\tilde{O}(T + \alpha + \beta)$ rounds

Graph Family	Quality of LCS
General Graphs	$\alpha + \beta = O(D + \sqrt{n})$
Planar Graphs	$\alpha + \beta = \tilde{O}(D)$
Expander Graphs	$2^{\sqrt{\log n}}$
Graph with $D = O(1)$	$\alpha + \beta = \tilde{O}(n^{\frac{D-2}{2D-2}})$

Approximate min-cut

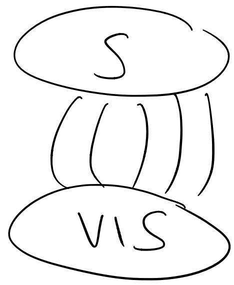

cut is a graph partitioning into $(S, V\setminus S)$

the value of the cut $(S, V\setminus S)$ is

$$\delta_G(S) = |\{(u,v) \mid u \in S,\ v \in V\setminus S\}|$$

min-cut is $\min_{S \subseteq V} \delta_G(S)$

O(log n) Approximation

Karger's edge sampling: Given graph $G = (V, E)$ with min cut λ.

Let $G[p]$ be the subgraph obtained by sampling each edge $e \in G$ indep. w.p.

$p = \frac{100 \log n}{\varepsilon^2 \lambda}$. Then: w.h.p. for $S \subseteq V$:

$$\delta_{G[p]}(S) \in (1 \pm \varepsilon)\, \delta_G(S) \cdot p$$

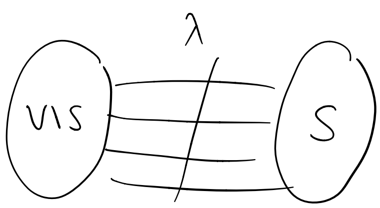

Cor: $G[p]$ for $p = \Theta\left(\frac{\log n}{\lambda}\right)$ then w.h.p. $G[p]$ is connected.

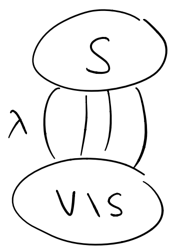

- on the one hand: $G[p_1]$ is connected w.h.p. for $p_1 = \Theta\left(\frac{\log n}{\lambda}\right)$.
- on the other hand: $G[p_2]$ is connected w.p. $\leq \frac{3}{4}$ for $p_2 = \frac{1}{\lambda}$.

Algorithm

$O(\log^2 n)$ Sampling experiments

For $i = 1$ to $\lceil \log n \rceil$

For $j = 1$ to $\ell = O(\log n)$

$G_{ji} = G[p_j]$ for $p = 2^{-j}$

$X_{ji} = 1$ iff G_{ji} is connected. → Connectivity can be solved in $\tilde{O}(D + \sqrt{n})$.

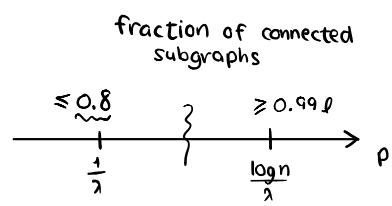

$j^* = \max j$ s.t. $\sum_{i=1}^{\ell} X_{ji} \geq 0.9 \cdot \ell$

Return 2^{j^*}

FOCS 2020: Exact min-cut in $\tilde{O}(D+\sqrt{n})$. $(1+\varepsilon)$ approximation with shortcuts.

LB: for unweighted min-cut: $\Omega(\lambda \cdot D)$ $\hookrightarrow$ Poly(D)

K-sparse certificate

$H \subseteq G$ is a K-sparse certificate if:

1) $\forall S \subseteq V$, $\delta_H(S) \geq \min\{K, \delta_G(S)\}$

2) $|E(H)| \leq K \cdot n$.

Centralized Construction

$H \leftarrow \emptyset$, $G' \leftarrow G$

For $i=1$ to k:

F_i = MSF in G'

$H \leftarrow H \cup F_i$

$G' \leftarrow G' \setminus F_i$

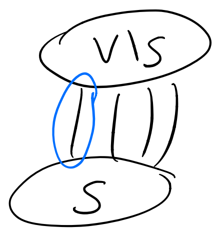

Correctness:

Focus on a specific cut $(S, V \setminus S)$

In every step i: collect at least one cut edge (if exists).

Distributed Implementation

Phase i: edges in $F_1 \cup \cdots \cup F_{i-1}$ have weight ∞

edges in $G \setminus \bigcup_{j=1}^{i-1} F_j$ have weight 1

Compute MST.

Lemma: Can compute K-certificate in $\tilde{O}(K(D+\sqrt{n}))$ rounds.

Next steps: - assume λ is known.

- provide alg' that works in $\tilde{O}(\lambda(D+\sqrt{n}))$ rounds.

High Level Idea of Centralized Construction (Matula's)

- Compute K'-certificate $H \subseteq G$ for $K' = (1+\frac{\varepsilon}{10})\lambda$

  $|E(H)| \leq (1+\frac{\varepsilon}{10})\lambda n$
  Every min-cut in H is a min-cut in G and vice versa.

- Define G' obtained by contracting all edges in $G \setminus H$.

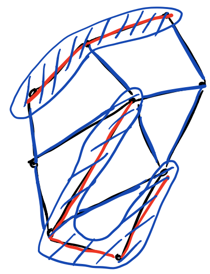

Case 1: $|V(G')| \geq \frac{n}{1+\frac{\varepsilon}{5}}$

Avg deg in $G' \leq \frac{2|E(G')|}{|V(G')|} \leq \frac{2(1+\frac{\varepsilon}{10})\lambda n}{\frac{n}{1+\frac{\varepsilon}{5}}} \leq (2+\varepsilon)\lambda$

$\Rightarrow$ $\exists$ node (possibly super-node) with deg $\leq (2+\varepsilon)\lambda$ $\Rightarrow$ defines a cut.

Case 2: $|V(G)| \leq \frac{n}{1+\frac{\varepsilon}{5}}$

$\Rightarrow$ continue recursively.

Distributed Implementation

$G \setminus H$ to be contracted:

Every super-node in G' corresponds to a connected comp. of H.

Compute certificate for G'.

- assign weight of 0 to all edges inside super-node.

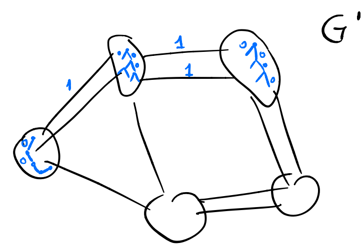

Complexity: $\tilde{O}(\lambda(D+\sqrt{n}))$ rounds.

Goal: $(2+\varepsilon)$-approx in $\tilde{O}(D+\sqrt{n})$ rounds.

Idea: Use Karger to sparsify the graph. $G' = G[p]$ for $p = \frac{200 \log n}{\varepsilon^2 \lambda}$:

$$\forall S \subseteq V, \quad \delta_{G'}(S) \in (1 \pm \varepsilon)\, \delta_G(S) \cdot p$$

$\Rightarrow$ Apply Matula's alg on G' in $\tilde{O}(D+\sqrt{n})$ rounds.

Omit assumption on λ

Try $\lambda = (1+\frac{\varepsilon}{10})^i$ as a guess for $i = 1, 2, 3, \ldots$