All-or-Nothing multicommodity flow and related problems

PhD Research Proposal

Arindam Pal
Department of Computer Science and Engineering
Indian Institute of Technology, Delhi

Multicommodity Flow

- Input: Graph $G = (V, E)$ with n nodes and m edges with edge (node) capacity.
- Set of terminal pairs $T = \{(s_1, t_1), \ldots, (s_k, t_k)\}$.
- Demand of a pair $(s_i, t_i) = d_i$.
- Objective: Satisfy the demands of maximum number of (s_i, t_i) pairs while respecting the capacity constraints.
- We will focus on edge problems in general graphs.

PROBLEM VARIATIONS

- Edge-disjoint paths (EDP): For each pair (s_i,t_i), $d_i = 1$ and there are edge-disjoint paths P_i joining s_i and t_i.
- Unsplittable flow (UFP): For each pair (s_i,t_i), d_i flow can be routed on a path P_i joining s_i and t_i.
- All-or-Nothing multicommodity flow (ANMF): There is a multicommodity flow f such that d_i flow is routed between s_i and t_i.
- Minimum multicut (MULTICUT): A set of edges C whose removal separates each pair (s_i,t_i). Goal is to find a multicut C of minimum capacity.

Approximability of ANMF

- NP-hard and APX-hard, even on trees.
- In trees, reduces to the maximum integer multicommodity flow problem.
- 2-approximation known for the unweighted case and 4-approximation for the weighted case.
- For general graphs, the best approximation factor is $O(\log^2 k)$ and for planar graphs, it is $O(\log k)$.

Problems we want to work on

- Improving the best known approximation factor $O(\log^2 k)$ for ANMF in general graphs.
- To give a $O(\beta^c)$-approximation algorithm for ANMF, where β is the *maxflow-mincut gap* for the graph and c is a constant. For general graphs $\beta = O(\log k)$, and for planar graphs $\beta = O(1)$.
- To improve the approximation factors for UFP on line and ring graphs (78.51 and 79.51 + ε) to small constants.
- Closing the gap between upper and lower bounds for EDP and MULTICUT.

LP for Multicommodity Flow

$\mathcal{P}_i$: set of paths joining s_i and t_i.
$\mathcal{P} = \bigcup_i \mathcal{P}_i$.
$f(P)$: amount of flow sent on path $P \in \mathcal{P}$.
x_i: total flow sent on all paths in $\mathcal{P}_i$.

$$\max \quad \sum_{i=1}^{k} x_i$$

$$\text{such that} \quad x_i = \sum_{P \in \mathcal{P}_i} f(P) \qquad 1 \le i \le k$$

$$\sum_{P: e \in P} f(P) \le c_e \qquad \forall e \in E$$

$$x_i, f(P) \in [0,1] \qquad 1 \le i \le k, P \in \mathcal{P}$$

Properties of the LP

- The number of fractionally routed paths in a basic solution to the LP is at most m.
- Hence, we can assume $c_e \leq m$ for all $e \in E$.
- By making parallel copies of edges, we can assume that G has only unit capacity edges.
- We can work with multicommodity flows where the flow on each path is $\Omega(1/\log n)$.

Flow decomposition lemma

Let $\bar{f}$ be a feasible solution to a multicommodity flow instance. Then there is a solution $\bar{f}'$ for the same instance such that

(i) $|\bar{f}'| = \Omega(|\bar{f}|)$,

(ii) there is a flow decomposition for $\bar{f}'$ such that the flow on each flow path is at least $\frac{c}{\log n}$ for some constant c.

By duplicating terminals we can assume that the flow for each pair is $\Theta(\frac{1}{\log n})$ and that the number of pairs is $\Theta(|f| \log n)$.

Proof of FD lemma

$$l = \lceil f_i \log n \rceil$$

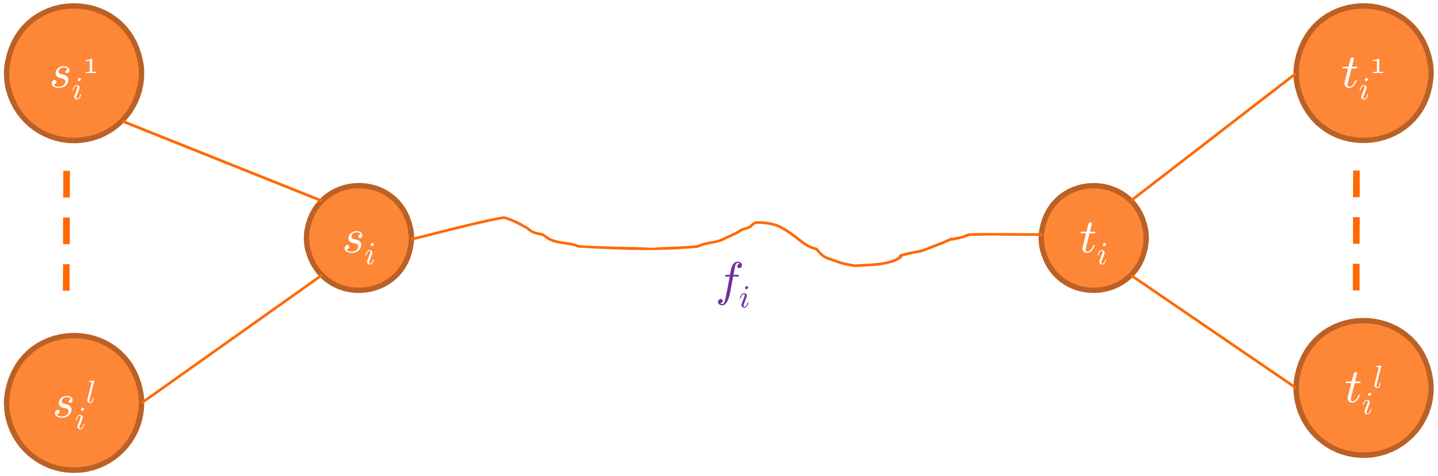

Proof continued…

- Flow for s_i^j and $t_i^j = 1/\log n$, if $j < l$.
- Flow for s_i^l and $t_i^l = f_i - (l-1)/\log n$.
- Round the flow to $1/\log n$ with probability $f(P)\,\log n$ and to 0 otherwise.
- Expected total flow = $f(P)$.
- For a sufficiently large constant c, the probability that the flow on any edge exceeds c is small (Chernoff bound).

PROOF CONTINUED…

- Set the flow on each path to $1/(c \log n)$.
- None of the (unit) edge capacities are violated.
- Total flow $= \Omega(|f|)$.
- Flow for each pair is $1/(c \log n)$ and each pair uses a single path.
- Number of pairs $= \Theta(|f| \log n)$.

Product MC Flow and Sparsest Cut

- For each pair of nodes (u, v), $d(u, v) = w(u)w(v)$, where w: $V \to \mathrm{R}^+$ is a weight function on nodes.
- Maximum concurrent flow is the largest λ such that λd flow can be routed in G.
- Sparsity of a cut is the ratio of the capacity of the cut and the demand separated by the cut.
- The worst case (maximum) ratio of minimum sparsity to maximum concurrent flow is the maxflow-mincut gap (β).
- For product MCF, $\beta = O(\log k)$ in general graphs and $\beta = O(1)$ in planar graphs.

Flow well-linked sets

- π: $X \to R^+$ is a weight function on a set of nodes X.
- Demand $d(u, v) = \pi(u)\pi(v)/\pi(X)$.
- X is π–flow–linked if there is a feasible multicommodity flow with demand $d(u, v)$.
- If $\pi(v) = \alpha$ for all $v \in X$, then X is α–flow–linked.
- In this case, $d(u, v) = \alpha/|X|$.

Cut well-linked sets

- X is cut-well-linked in G if for every cut $(S, V-S)$, the number of edges cut is at least the number of X vertices on the smaller side.

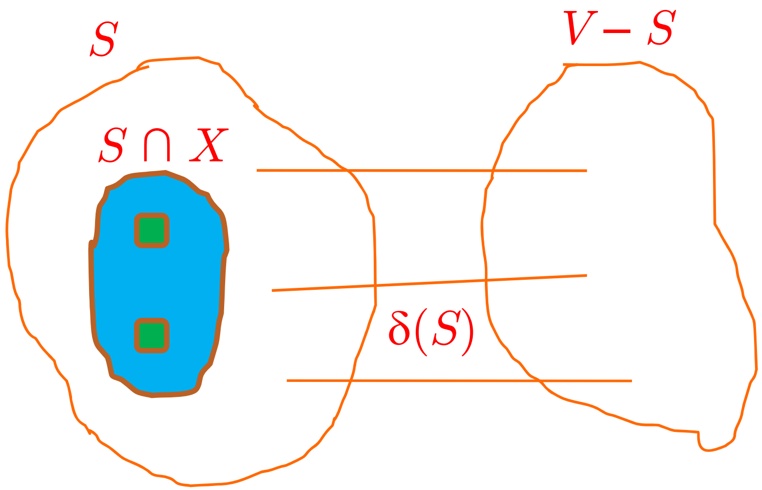

- $|\delta(S)| \geq |S \cap X|$ for all S such that $|S \cap X| \leq |X|/2$.

Cut well-linked sets definitions

- π: $X \to \mathrm{R}^+$ is a weight function on a set of nodes X.
- X is π–edge–cut–linked if $|\delta(S)| \geq \pi(S \cap X)$ for all S such that $\pi(S \cap X) \leq \pi(X)/2$.
- Equivalently, $|\delta(S)| \geq \min\{\pi(S \cap X), \pi(S^c \cap X)\}$.
- X is π–node–cut–linked if $|N(S)| \geq \pi(S \cap X)$ for all S such that $\pi(S \cap X) \leq \pi(X)/2$.
- $N(S)$ is the neighborhood of S.
- In the example, $\pi(X) = |X|$.

Relation among cut and flow linked

- X is π–flow–linked $\Rightarrow$ X is π/2–cut–linked.
- X is π–cut–linked $\Rightarrow$ X is π/β–flow–linked.

α–FLOW–LINKED $\Rightarrow$ $\alpha/2$–CUT–LINKED

$d(u, v) = \alpha/|X|$ $\quad$ $r = |S \cap X|$ $\quad$ αr <= $\alpha|X|/2$

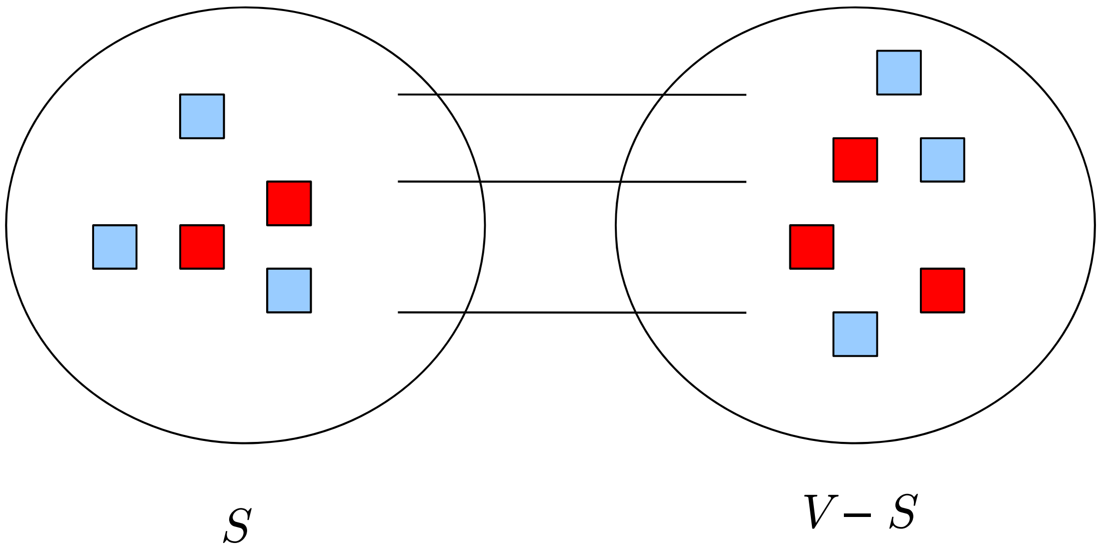

$|\delta(S)| \geq \alpha r(|X| - r) \,/\, |X|$ $\quad$ $|\delta(S)| \geq \alpha r/2$

π–FLOW–LINKED $\Rightarrow \pi/2$–CUT–LINKED

- Given: X is π–flow–linked.

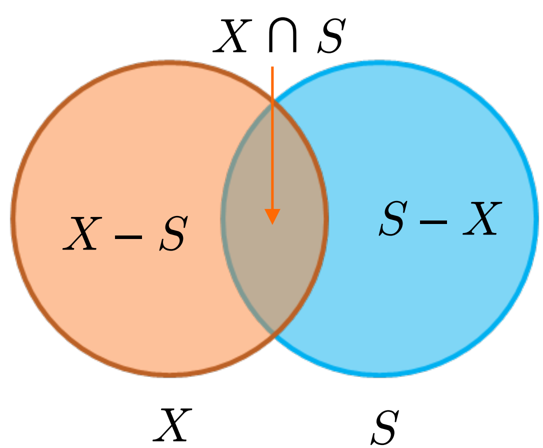

- Required to prove: $|\delta(S)| \geq \pi(S \cap X)/2$ for all S such that $\pi(S \cap X) \leq \pi(X)/2$.

PROOF

Since, $\pi(X \cap S) \leq \frac{\pi(X)}{2}$, therefore $\pi(X - S) \geq \frac{\pi(X)}{2}$.

$$
\begin{aligned}
|\delta(S)| &\geq |\delta(X \cap S, X - S)| \\
&= \sum_{\substack{u \in X - S \\ v \in X \cap S}} d(u, v) \\
&= \sum_{u,v} \frac{\pi(u)\pi(v)}{\pi(X)} \\
&= \frac{1}{\pi(X)} \sum_{u \in X - S} \pi(u) \sum_{v \in X \cap S} \pi(v) \\
&= \frac{\pi(X - S)\pi(X \cap S)}{\pi(X)} \\
&\geq \frac{1}{2}\pi(X \cap S).
\end{aligned}
$$

Proposition

- Given a subset $X \subseteq V$ of nodes that is α–flow–linked in G, any matching M on X can be fractionally routed with congestion 2/α. Hence, $\Omega(\alpha|X|)$ flow can be routed for M.

Proof

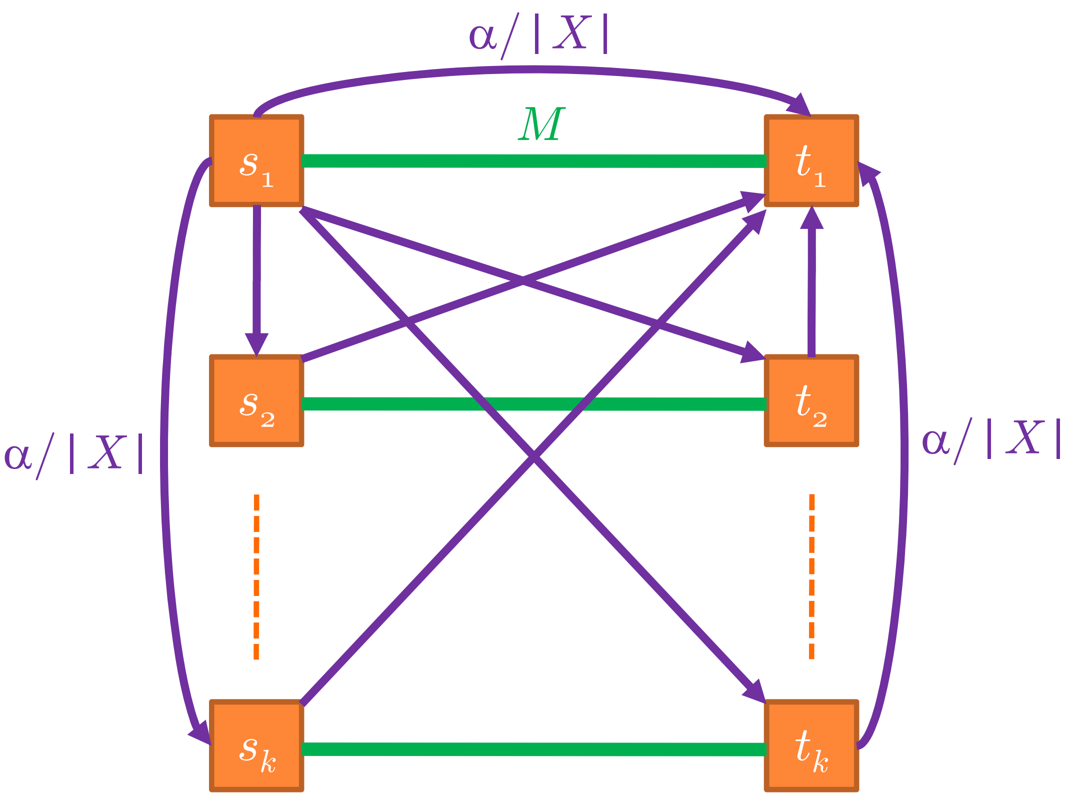

Proof continued…

- There is a feasible MC Flow with demand $\alpha/|X|$ between any two nodes in X.
- Unit demand between any two nodes in M.
- Route flow using intermediate nodes.
- Number of intermediate nodes used $N = |X|/\alpha$.
- Congestion $= N/|X| = 1/\alpha$.

Well-linked Sets

- Input instance: G, X, M.

- X: set of terminals $= \{s_1, t_1, \ldots, s_k, t_k\}$.
- M: matching on $X = (s_1, t_1), \ldots, (s_k, t_k)$.

- X is π-well-linked in G for some π: $X \rightarrow \mathrm{R}^+$.

Examples of Well-linked Sets

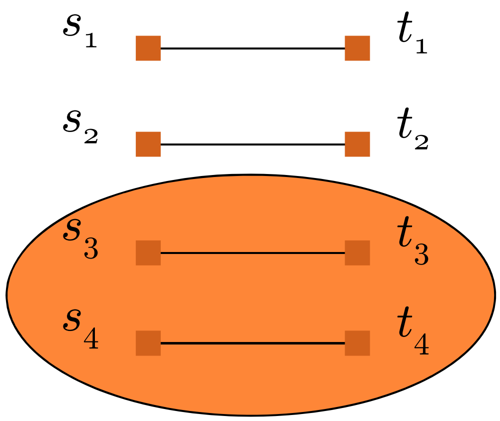

NOT a well-linked instance

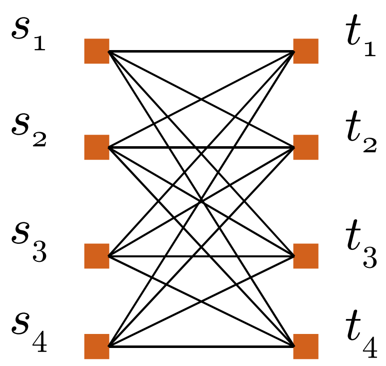

A well-linked instance

Well-linked Decomposition

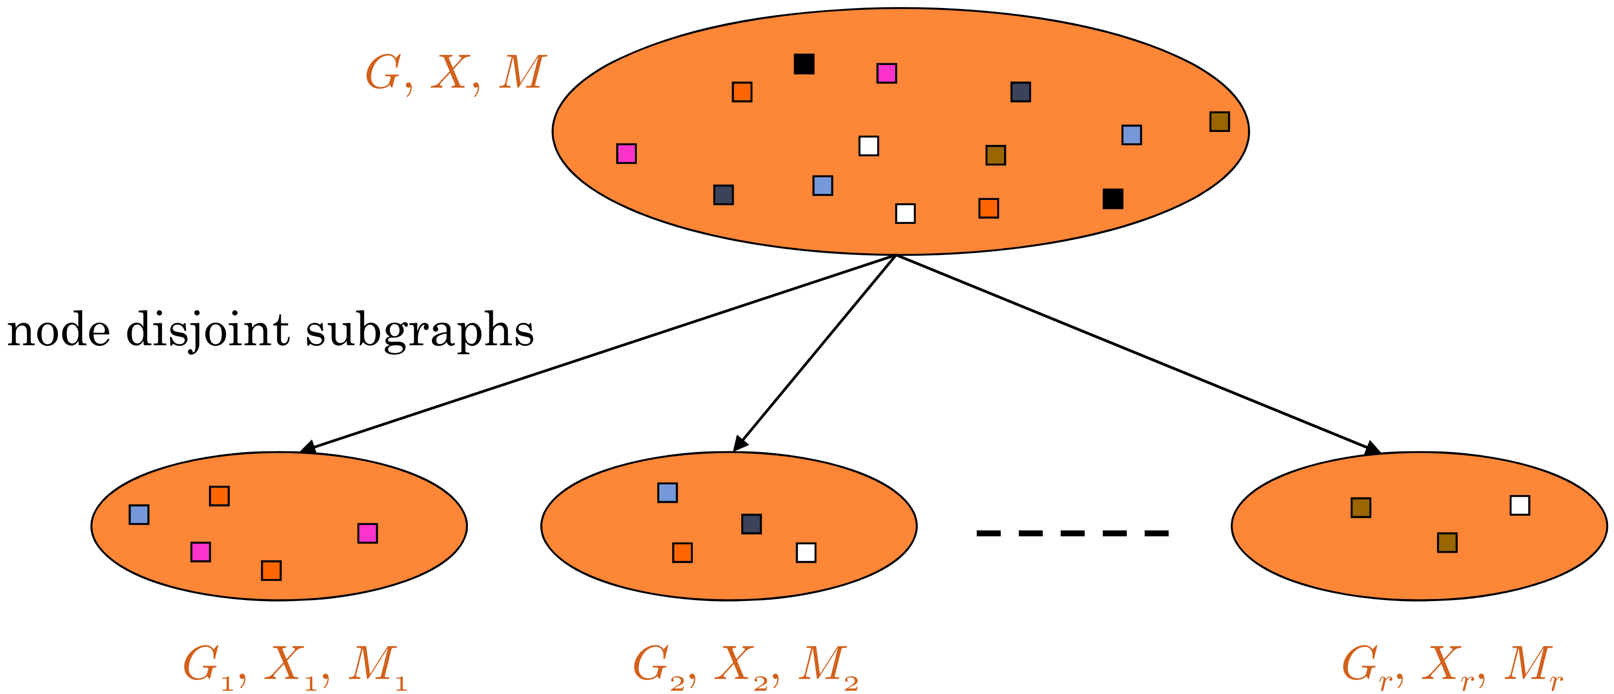

$M_i \subseteq M$

X_i is well-linked in G_i

$\sum_i |X_i| \geq \mathrm{OPT}/\alpha$

Example of Well-linked Decomposition

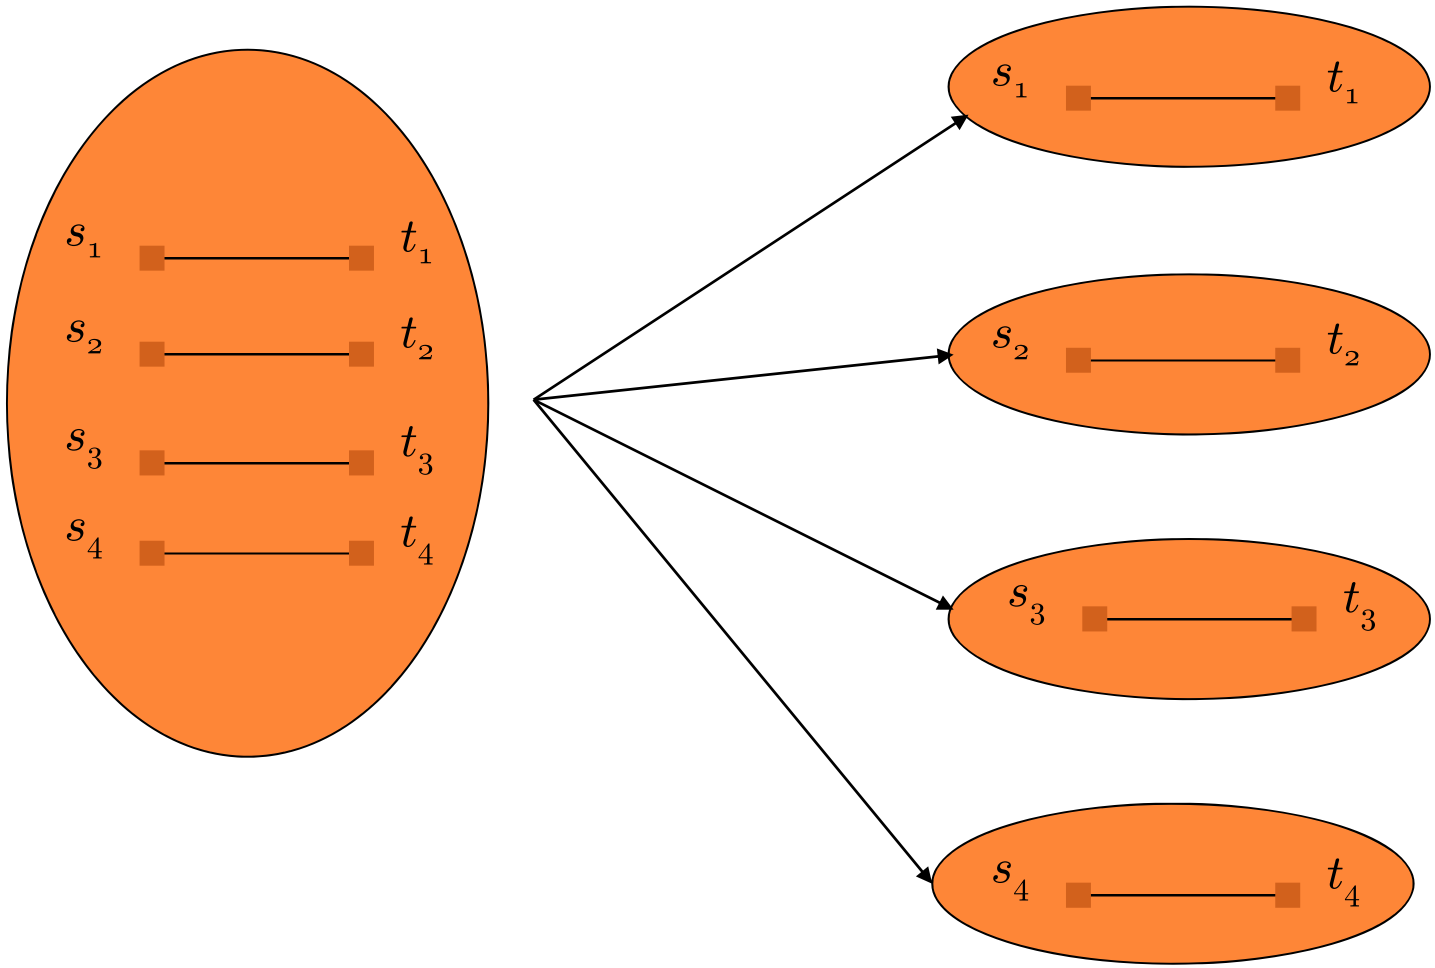

Decomposition using Sparse Cuts

- Start with a MC flow f for X in G with flow f_i for the pair s_i, t_i.
- f_i is decomposed into flow paths.
- γ(H) is the total flow induced in subgraph H by the original flow f.
- γ(u, H) is the flow in H for u.
- OPT is a solution to the LP for EDP.
- X is the set of terminals (sources and sinks).

Decomposition Algorithm

- Suppose $\gamma(H) < \beta \log \mathrm{OPT}$. Let u, v be some pair with positive flow in H. Define $\pi(u) = \pi(v) = 1$ and $\pi(w) = 0$ for $w \neq u, v$. Stop and output (H, π).
- If X is $\pi/(10\, \beta \log k)$-flow-linked, stop and output (H, π) $[\pi(u) = \gamma(u, H)/(10\, \beta \log \mathrm{OPT})]$.
- Else

  Find a sparse cut $(S, V - S)$ w.r.t. π in G.

  Remove flow on edges of $\delta_G(S)$.

  $G_1 = G[S]$, $G_2 = G[V - S]$.

  Recurse on G_1 and G_2 with remaining flow.

Analysis

- Subgraphs at the end of recursion:
  (G_1, X_1, π_1) , ... , (G_l, X_l, π_l).
- π_i is the remaining flow for X_i.
- X_i is $\pi_i/(10\ \beta \log k)$-flow-linked in G_i.
- $\sum_i \pi_i(X_i) \geq$ Original flow – No. of edges cut.

BOUNDING THE NUMBER OF EDGES CUT

- X is not $\pi/(10\ \beta \log k)$-flow-linked

$$\Rightarrow |\delta_G(S)| \leq \pi(S)/(10 \log k).$$

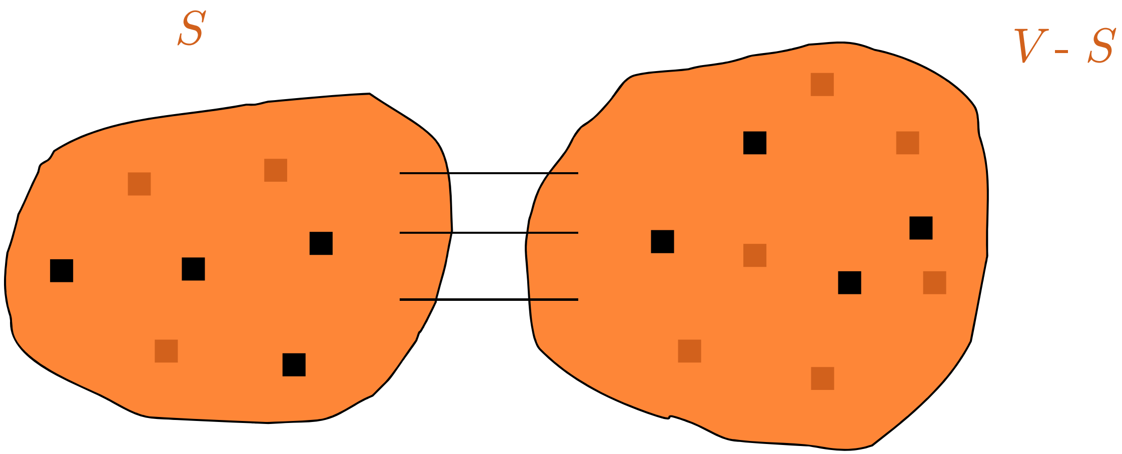

Analysis continued…

- Lemma: Total number of edges cut $\leq f/2$.
- $T(f)$: max no. of edges cut if started with flow f.

$$T(f) \leq T(f_1) + T(f_2) + f_1 / 10 \log k$$

- For $f \leq k$, $T(f) \leq f/2$.
- $\sum_i \pi_i(X_i) \geq f/2$.
- X_i is $\pi_i/(10\,\beta \log k)$-flow-linked in G_i.

Lower bounds

- Theorem: ANMF is at least as hard to approximate as MAXIMUM INDEPENDENT SET. Hence, it is impossible to approximate ANMF to within a factor of $\Omega(m^{1/2-\varepsilon})$, for any fixed $\varepsilon > 0$, unless P = NP.

- Proof: Using a polynomial time approximation factor preserving reduction from MAXIMUM INDEPENDENT SET to ANMF.

Some ideas

- Select at most one path $p_i \in P_i$ with probability f_{p_i} $(i = 1, \dots, k)$, obtained by solving MCF-LP. Let $p_1, \dots, p_k$ be the set of paths chosen.
- Let $(p_{\pi(1)}, \dots, p_{\pi(k)})$ be the permutation of $(p_1, \dots, p_k)$ such that $d_{\pi(1)} \geq d_{\pi(2)} \geq \dots \geq d_{\pi(k)}$.
- Add the path $p_{\pi(i)}$ if adding it does not violate the capacity constraint for every edge $e \in p_{\pi(i)}$.
- If $p_{\pi(i)}$ is added, the demand $d_{\pi(i)}$ is routed along the path $p_{\pi(i)}$.
- We can also consider a set of paths $\mathbf{P}_i \in P_i$ in decreasing order of demands by either sorting all
- paths in $\mathbf{P}_1 \cup \dots \cup \mathbf{P}_k$ or sorting each $\mathbf{P}_i$ separately.

References

- [CCGK07] A. Chakrabarti, C. Chekuri, A. Gupta, A. Kumar. Approximation Algorithms for the Unsplittable Flow Problem. *Algorithmica* 47(1): 53-78 (2007).
- [CKS04] C. Chekuri, S. Khanna, F. B. Shepherd. The All-or-Nothing Multicommodity Flow Problem. *ACM Symposium on Theory of Computing*, June 2004.
- [CKS05] C. Chekuri, S. Khanna, F. B. Shepherd. Multicommodity Flow, Well-linked Terminals, and Routing Problems. *ACM Symposium on Theory of Computing*, May 2005.

REFERENCES

- [GKRSM99] V. Guruswani, S. Khanna, R. Rajaraman, B. Shepherd, M. Yannakakis. Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems. *ACM STOC* 19-28 (1999).
- [GVY96] N. Garg, V. V. Vazirani, M. Yannakakis. Approximate max-flow min-(multi)cut theorems and their applications. *SIAM Journal on Computing* 25(2): 235-251 (1996).
- [GVY97] N. Garg, V. V. Vazirani, M. Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica* 18:3-20 (1997).

References

- [K96] J. Kleinberg. Approximation Algorithms for Disjoint Paths Problems. *PhD thesis*, MIT (1996).
- [LR99] T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM* 46(6):787-832 (1999).
- [R08] H. Racke. Optimal Hierarchical Decompositions for Congestion Minimization in Networks. *ACM Symposium on Theory of Computing*, May 2008.

Questions?

Thank You!