Fair Network Design via Iterative Rounding

Aditi Laddha, Mohit Singh, and Santosh Vempala

Georgia Institute of Technology

Introduction

Survivable Network Design Problem

Given an undirected multigraph $G = (V, E)$ with non-negative edge costs and an integral cut requirement function $f: 2^V \to \mathbb{Z}_{\geq 0}$, the goal is to find a subgraph of G with the minimum cost that satisfies the cut requirements.

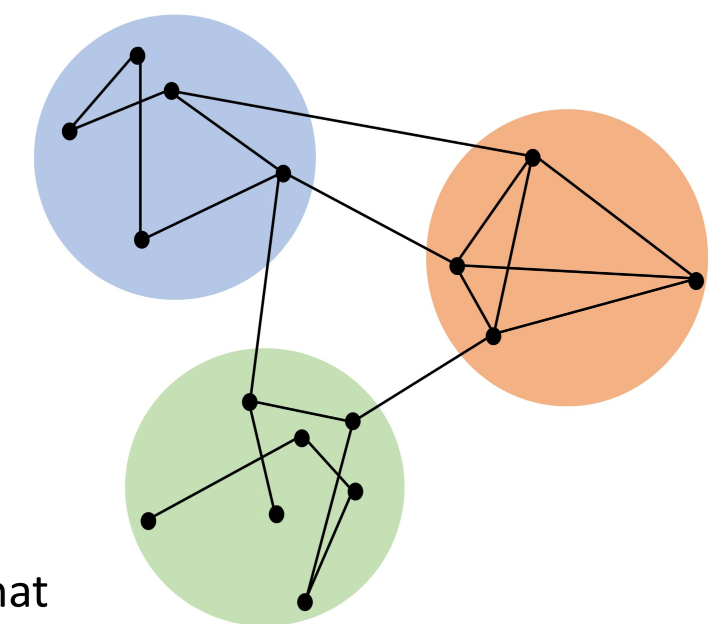

Generalized Steiner Network Problem

Special subcase of SNDP where the cut requirement function can be expressed as a pairwise connectivity function $r: V \times V \to \mathbb{Z}_{\geq 0}$, and the goal is to find a minimum-cost subgraph of G that contains at least r_{ij} edge-disjoint paths for each pair $(i, j) \in V \times V$.

This problem captures:

- Minimum Steiner trees
- Minimum Steiner forests
- Minimum k-edge connected subgraph

Theorem.[Jain[1]] The survivable network design problem with a skew supermodular requirement function can be approximated to within a factor of 2 in polynomial time, given access to an oracle that solves its LP relaxation.

Multi-Criteria

What if given a network and certain connectivity requirements, multiple agents with different (and possibly conflicting) cost functions have to agree on a common architecture? We want to be 'fair' to all the agents.

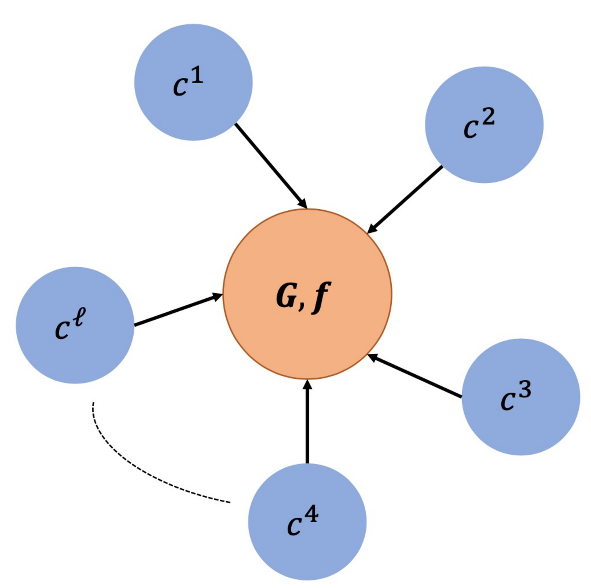

Fairness? Many options like Pareto Optimality[4] and Multi-Objective Approximation[5] are studied in literature. We focus on minimizing the maximum cost among all the agents:

Max-linear Optimization

Given a feasible set $\mathcal{F} \subseteq 2^E$ and $\ell \geq 2$ linear cost functions $c^1, c^2, \ldots, c^\ell$ with $c^i: \mathcal{F} \to \mathbb{R}_+$, the goal of Max-linear optimization problem is to find

$$\min_{X \in \mathcal{F}} \max_{i \in [\ell]} c^i(X).$$

Results

We use the following LP relaxation for the Multi-Criteria Survivable Network Design problem (MCSNDP).

$LP\ (G, f, \{c^i\}_{i=1}^{\ell})$

$$\begin{aligned} &\min \quad z \\ \text{s.t.} \quad &\sum_{e \in \delta_G(S)} x_e \geq f(S) \quad \forall S \subseteq V \\ &\sum_{e \in E(G)} c_e^i x_e \leq z \quad \forall i \in \{1, \ldots, \ell\} \\ &z \geq 0 \\ &1 \geq x_e \geq 0 \quad \forall e \in E(G). \end{aligned}$$

Skew Supermodular

A function $f: 2^V \to \mathbb{Z}$ is called skew supermodular if for any $A, B \subseteq V$, at least one of the following holds:

- $f(A) + f(B) \leq f(A \cap B) + f(A \cup B)$
- $f(A) + f(B) \leq f(A \setminus B) + f(B \setminus A)$.

$f(S) = \max_{i \in S, j \notin S} r_{ij}$ is skew supermodular.

Due to skew supermodularity and integrality of f, we get certain nice properties at extreme points.

Extreme Point Support:

Let x be an extreme point solution to $LP\ (G, f, \{c^i\}_{i=1}^{\ell})$ with $0 < x_e < 1$ for each edge $e \in E$. Then there exists a laminar family, B, of tight sets satisfying the following:

- $|E| - (\ell - 1) \leq |B| \leq |E| + 1$, and
- The vectors $\{\chi(\delta(S)), S \in B\}$ are linearly independent.

Theorem. For $\ell \geq 2$ and any skew supermodular integral function f, let x be an extreme point solution to $LP\ (G, f, \{c^i\}_{i=1}^{\ell})$. Then there exists an edge e with $x_e \geq 1/\ell$.

Remark: This generalizes Jain's result.

Iterative Multi-Criteria Network Design Algorithm Outline

Input: A graph G, a skew supermodular function f, and a set of $\ell \geq 2$ cost functions $\{c^i\}_{i=1}^{\ell}$.

1. Initialize $F = \emptyset$.
2. While $f \neq 0$:
   1. Find an optimal extreme point solution x to $LP\ (G, f, \{c^i\}_{i=1}^{\ell})$.
   2. Add all edges with $x_e \geq 1/\ell$ to F.
   3. Update G and f.

Output: F.

Integrality Gap:

The integrality gap of $LP\ (G, f, \{c^i\}_{i=1}^{\ell})$ is ℓ.

For $\ell \geq 2$, consider a multi-graph containing 2 nodes u, v with ℓ edges $\{e_1, e_2, \ldots, e_\ell\}$.

- Demand function is $f(\{u\}) = f(\{v\}) = 1$.
- Cost functions:

$$c_{e_j}^i = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}.$$

- The integral optimal value is 1 but $z_{LP} = 1/\ell$.

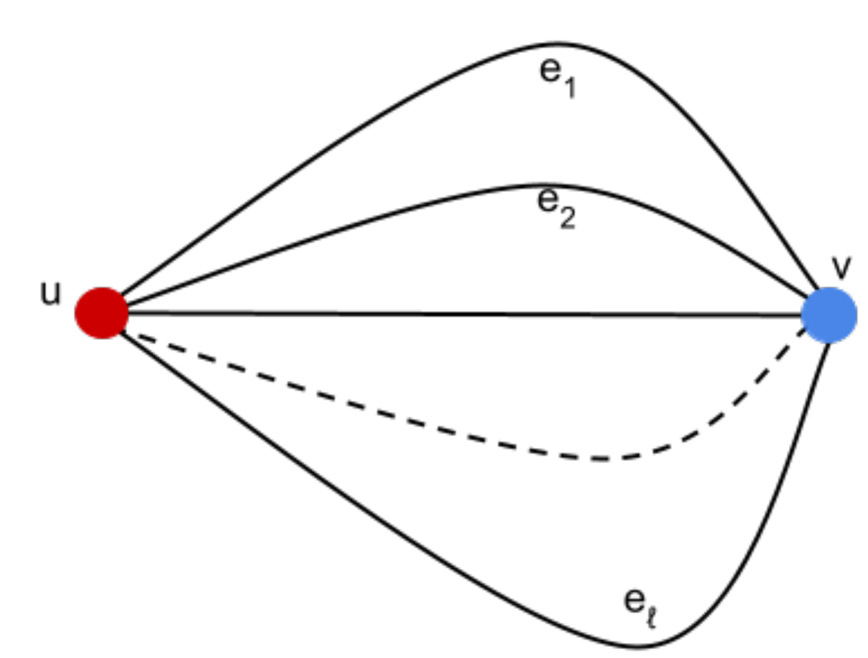

Conclusion

Theorem. For $\ell \geq 2$, there is a polynomial-time ℓ-approximation algorithm for the Fair generalized Steiner network problem with ℓ cost functions.

Remarks.

- The algorithm and the guarantee extend to the multi-criteria survivable network design problem with a skew supermodular cut requirement function provided the LP relaxation can be solved in polynomial time.
- Fair network design for 2 players has no penalty in the approximation factor!

Hardness of Approximation

As a consequence of the hardness of approximation of multi-criteria shortest path[3], we get the following hardness result.

Theorem. The multi-criteria generalized Steiner network problem with ℓ cost functions is not approximable within $\log^{1-\epsilon} \ell$ for any $\epsilon > 0$, unless $\mathrm{NP} \subseteq \mathrm{DTIME}(n^{\mathrm{poly}(\log n)})$.

Future Work

1. Can we get better approximation algorithms using stronger LP relaxations?
2. Can we improve the hardness of approximation factor to match the upper bound?
3. Can we get better approximation algorithms for multi-criteria versions of special subclasses of network design problems, e.g., uniform k-connectivity?

References

1. Kamal Jain. A factor 2 approximation algorithm for the generalized steiner network problem. Combinatorica, 21(1):39–60, 2001.
2. Adam Kasperski and Paweł Zieliński. On the approximability of minmax (regret) network optimizationproblems.Information Processing Letters, 109(5):262–266, 2009.
3. Viswanath Nagarajan, R Ravi, and Mohit Singh. Simpler analysis of lp extreme points for traveling salesman and survivable network design problems. Operations Research Letters, 38(3):156–160, 2010
4. Christos H Papadimitriou and Mihalis Yannakakis. On the approximability of trade-offs and optimal access of web sources. In Proceedings 41st Annual Symposium on Foundations of Computer Science, pages 86–92. IEEE, 2000.
5. Ramamoorthi Ravi, Madhav V Marathe, Sekharipuram S Ravi, Daniel J Rosenkrantz, and Harry B Hunt III. Many birds with one stone: Multi-objective approximation algorithms. In Proceedings of the twenty-fifth annual ACM symposium on Theory of computing, pages 438–447, 1993.
6. Uthaipon Tantipongpipat, Samira Samadi, Mohit Singh, Jamie H Morgenstern, and Santosh S Vempala. Multi-criteria dimensionality reduction with applications to fairness. Advances in neural information processing systems, (32), 2019.