



Socially fair network design via iterative rounding

Aditi Laddha^{*}, Mohit Singh, Santosh S. Vempala

Georgia Institute of Technology, North Ave NW, 30332, Atlanta, USA

ARTICLE INFO

Article history:

Received 18 November 2021

Received in revised form 11 July 2022

Accepted 27 July 2022

Available online 3 August 2022

Keywords:

Approximation algorithm

Iterative rounding

Network design

ABSTRACT

We show that the *multi-criteria* generalized Steiner network problem with $\ell \geq 2$ cost functions has a polynomial-time ℓ -approximation algorithm. This generalizes the celebrated result of Jain for a single cost function. We also show that the approximation factor is tight for the natural LP relaxation. The problem is motivated by network design where multiple players have different or competing valuations and must agree on a common network (e.g., infrastructure design).

© 2022 Elsevier B.V. All rights reserved.

1. Introduction

Given an undirected multigraph $G = (V, E)$ with non-negative edge costs and a cut requirement function $f : 2^V \rightarrow \mathbb{Z}_{\geq 0}$, the *Survivable Network Design Problem* asks to find a subgraph of G with the minimum cost that satisfies the cut requirements. A widely studied special case of this problem is the *generalized Steiner network problem* where we are given a pairwise connectivity function $r : V \times V \rightarrow \mathbb{Z}_{\geq 0}$ 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 several classical and widely studied connectivity problems like the minimum Steiner tree, minimum k -edge connected subgraph, etc. While even these special cases are NP-hard, several techniques have been developed to successfully design approximation algorithms for them. The culmination of this line of work was Jain's [11] elegant 2-factor approximation for the setting when the cut requirement function is *skew supermodular*, a condition that includes the generalized Steiner network problem. In his paper, Jain introduced the technique of iterative rounding of linear programming relaxations for combinatorial optimization, which has since been applied to many other problems.

In this paper, we consider the *multi-criteria* Survivable network design problem (MCSNDP). In this generalization, we are given an undirected multigraph, $G = (V, E)$, and cut requirements (as before), along with ℓ non-negative cost functions on the edges, $c^1, c^2, \dots, c^\ell : E \rightarrow \mathbb{R}_+$. We wish to find a subgraph of G that minimizes the maximum, over the ℓ cost functions of a subgraph that

has the required number of edges crossing every cut. The problem is motivated by network design where multiple players have different or competing valuations and must agree on a common network (e.g., infrastructure design). In this setting, one candidate for a *fair* solution is to minimize the maximum cost to one of the players.

Related work The appearance of multiple and possibly conflicting objective functions is a common phenomenon, and many real-world optimization problems involve finding a solution that finds a trade-off between these objectives. Multi-criteria optimization is a widely studied area of Operations Research, Economics, and Computer Science [6,8,17]. Several different techniques have been used to deal with multi-criteria optimization including Goal Programming [1], Pareto Optimality [4,16], Multi-objective Approximation Algorithms [9,17]. In this paper, we focus on the Max-linear approach to multi-criteria optimization. Given a combinatorial optimization problem and multiple linear cost functions, the objective is to find a feasible solution that minimizes the maximum over all the cost functions. Formally, given a feasible set $\mathcal{P}$ and $\ell \geq 2$ linear cost functions $c^1, c^2, \dots, c^\ell$ with $c^i : \mathcal{P} \rightarrow \mathbb{R}_+$ for all $i \in [\ell]$, the goal of Max-linear optimization problem is to find

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

Max-linear versions of the assignment problem and other combinatorial optimization problems occur in the modeling of printed circuit boards assembly [5,14]. The Max-linear version of the minimum spanning tree problem was proved to be NP-hard in [10]. Recently, multi-criteria formulations similar to the Max-linear objective were used in [18,20] to give algorithms for fair PCA, and in [2,3] for approximating simultaneous Max-Cut.

^{*} Corresponding author.

E-mail addresses: aladdha6@gatech.edu (A. Laddha), mohit.singh@isye.gatech.edu (M. Singh), vempala@gatech.edu (S.S. Vempala).

Before we proceed further on fair network design, let us recall the classical approach for the problem with one cost function.

Iterative rounding The network design problem specified by a graph $G = (V, E)$, cut requirement function $f : 2^V \rightarrow \mathbb{Z}$, and a cost function $c : E \rightarrow \mathbb{R}_+$ is solved by the following integer program and its natural linear programming relaxation:

$$\begin{aligned}
 & \text{IP} \\
 & \text{minimize } \sum_{e \in E(G)} c_e x_e \\
 & \text{subject to } \sum_{e \in \delta_G(S)} x_e \geq f(S), \quad \forall S \subseteq V \\
 & \quad x_e \in \{0, 1\}, \quad \forall e \in E(G) \\
 & \text{LP} \\
 & \text{minimize } \sum_{e \in E(G)} c_e x_e \\
 & \text{subject to } \sum_{e \in \delta_G(S)} x_e \geq f(S), \quad \forall S \subseteq V \\
 & \quad 1 \geq x_e \geq 0, \quad \forall e \in E(G)
 \end{aligned}$$

We can now state Jain's seminal result.

Theorem 1 (Jain [11]). *The 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.*

In particular, the LP relaxation for the generalized Steiner network problem is polytime solvable. The following structural theorem is at the core of the above result.

Theorem 2 (Jain [11]). *For any network design problem with a skew supermodular requirement function, any extreme point solution to its LP relaxation with $x_e > 0$ for each e contains at least one edge e with $x_e \geq 1/2$.*

This theorem is the foundation of iterative rounding – the algorithm includes the edge of weight at least half in its solution, adjusts the requirements, solves the residual LP (which again has an edge of weight at least $1/2$), and continues. Similar structural theorems have been discovered for other problems, notably the bounded-degree MST problem [19]. The book [13] provides a comprehensive discussion of iterative rounding methods for combinatorial optimization.

Multi-criteria (fair) network design When we have two (or more) cost functions, can we (approximately) minimize the maximum of the two costs? Would iterative rounding still work, and if so with what approximation guarantee?

This multi-criteria problem can be captured by the following integer program and the corresponding linear programming relaxation denoted by $LP_{mcsndp}(G, f, \{c^i\}_{i=1}^l)$:

$$\begin{aligned}
 & LP_{mcsndp}(G, f, \{c^i\}_{i=1}^l) \\
 & \text{minimize } z \\
 & \text{subject to } \sum_{e \in \delta_G(S)} x_e \geq f(S), \quad \forall S \subseteq V \\
 & \quad \sum_{e \in E(G)} c_e^i x_e \leq z, \quad \forall i \in \{1, \dots, l\} \\
 & \quad 1 \geq x_e \geq 0, \quad \forall e \in E(G)
 \end{aligned}$$




Fig. 1. Tight Example for LP_{mcsndp} .

$$\begin{aligned}
 & x_e \in \{0, 1\}, \quad \forall e \in E(G) \\
 & LP_{mcsndp}(G, f, \{c^i\}_{i=1}^l) \\
 & \text{minimize } z \\
 & \text{subject to } \sum_{e \in \delta_G(S)} x_e \geq f(S), \quad \forall S \subseteq V \\
 & \quad \sum_{e \in E(G)} c_e^i x_e \leq z, \quad \forall i \in \{1, \dots, l\} \\
 & \quad 1 \geq x_e \geq 0, \quad \forall e \in E(G)
 \end{aligned}$$

Our main result is the following.

Theorem 3. *For $\ell \geq 2$, there is a polynomial time ℓ -approximation algorithm for the Fair generalized Steiner network problem with ℓ cost functions. The algorithm and guarantee extend to multi-criteria network design problem with a skew supermodular cut requirement function provided the LP relaxation can be solved in polynomial time.*

Corollary 4. *For $\ell \geq 2$, given bounds $C_1, C_2, \dots, C_\ell$ in addition to $(G, f, \{c^i\}_{i=1}^\ell)$, there is a polynomial time algorithm that finds (if possible) a subgraph which satisfies the cut function requirements and for which the i -th cost function is at most $\ell \cdot C_i$.*

The algorithm is based on iterative rounding and the following structural result.

Theorem 5. *For $\ell \geq 2$ and any skew supermodular integral function f , let x be an extreme point solution to $LP_{mcsndp}(G, f, \{c^i\}_{i=1}^l)$. Then there exists an edge e with $x_e \geq 1/\ell$.*

Perhaps most surprising is that for two cost functions, the approximation factor is 2, the same as for a single cost function. In other words, fair network design for 2 players has no penalty in the approximation factor.

Tight example The following simple example (see Fig. 1) shows that ℓ is a tight bound for iterative rounding as well as for the integrality gap of the LP relaxation. For $\ell \geq 2$, consider a multi-graph containing 2 nodes u, v with ℓ edges $\{e_1, e_2, \dots, e_\ell\}$. The demand function is $f(\{u\}) = f(\{v\}) = 1$. For $i \in [\ell]$, the cost function $c^i : E \rightarrow \mathbb{R}_{\geq 0}$ is defined as

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

The optimal integral solution simply selects one of the edges with cost 1. However, the extreme point solution for LP_{mcsndp} is $x_{e_i} = 1/\ell$ for all $i \in \{1, \dots, \ell\}$ with $\max_{i \in [\ell]} \sum_e c_e^i x_e = 1/\ell$.

Hardness of approximation Consider the multi-criteria Shortest Path problem where given a graph $G = (V, E)$, 2 vertices $s, t \in V$, and ℓ non-negative cost functions on the edges, $c^1, c^2, \dots, c^\ell : E \rightarrow \mathbb{R}_+$, the goal is to find a path between s and t that minimizes the maximum of the ℓ cost functions over all paths between s and t .

In [12], they prove inapproximability results for multi-criteria versions of polynomially solvable network design problems, including the multi-criteria Shortest Path problem. We rephrase Theorem 1 from [12] in terms of multi-criteria network design.

Theorem 6 (Theorem 1, [12]). *The multi-criteria Shortest Path problem with ℓ cost functions is not approximable within $\log^{1-\epsilon} \ell$ for any $\epsilon > 0$, unless $\text{NP} \subseteq \text{DTIME}(n^{\text{poly}(\log n)})$.*

As a consequence of this theorem, we get the following result about the inapproximability of multi-criteria network design.

Theorem 7. *The multi-criteria generalized Steiner network problem with ℓ cost functions is not approximable within $\log^{1-\epsilon} \ell$ for any $\epsilon > 0$, unless $\text{NP} \subseteq \text{DTIME}(n^{\text{poly}(\log n)})$.*

A very interesting question for future work is whether the factor ℓ can be improved.

1.1. Approach

We use iterative rounding to construct a feasible integral solution for $\text{LP}_{\text{mcsndp}}(G, f, \{c^i\}_{i=1}^\ell)$ whose value is at most ℓ times worse than the optimal. Using the combinatorial properties of the set of tight cuts in an extreme point solution, we prove that any fractional extreme point solution contains an edge with value at least $1/\ell$. This is proved using a token counting technique with fractional tokens in Section 4.

2. Preliminaries

Definition 1 (Skew supermodular). A function $f : 2^V \rightarrow \mathbb{Z}$ is called skew supermodular if for any $A, B \subseteq V$, at least one of the following holds:

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

Any function defined by $f(S) = \max_{u \in S, v \notin S} \{r_{uv}\}$ is skew supermodular. This generalizes the minimum Steiner tree problem, the minimum Steiner forest problem, the minimum k -edge-connected problem.

Lemma 8 ([7]). *Let $G = (V, E)$ be a graph and $f : 2^V \rightarrow \mathbb{Z}$ be a skew supermodular function. Then for any subgraph of G , H , the function $f(S) - |\delta_H(S)|$ is also skew supermodular.*

Definition 2 (Laminar family). A set family $\mathcal{S} \in \{S_1, S_2, \dots\}$ is called laminar if for any $S_i, S_j \in \mathcal{S}$, $S_i \cap S_j = \emptyset$ or $S_i \subseteq S_j$ or $S_j \subseteq S_i$.

Lemma 9 (Rank Lemma [13]). *Let $P = \{Ax \geq b, x \geq 0\}$ be a polytope and let x be an extreme point solution of P such that $x_e > 0$ for each e . Then any maximal number of linearly independent tight constraints of the form $A_i x = b_i$ for some row i of A equals the number of variables.*

For a set $F \subseteq E$, let $\chi(F)$ denote the characteristic vector of F , i.e., $\chi(F) \in \mathbb{R}^{|E|}$ and has 1 corresponding to each edge $e \in F$, and 0 otherwise. As a consequence of the Rank lemma and Lemma 4.3 from [11], we get the following theorem.

Theorem 10 (Extreme point support). *Let x be an extreme point solution to $\text{LP}_{\text{mcsndp}}(G, f, \{c^i\}_{i=1}^\ell)$ with $0 < x_e < 1$ for each edge $e \in E$. Then there exists a laminar family, $\mathcal{B}$, of tight sets and a set of cost function constraints, $\mathcal{C}$ satisfying the following:*

- $|\mathcal{B}| + |\mathcal{C}| = |E| + 1$.
- The vectors $\{\chi(\delta(S)), S \in \mathcal{B}\}$ and $\{\chi(S), S \in \mathcal{C}\}$ are linearly independent.
- For every set $S \in \mathcal{B}$, $f(S) \geq 1$.

3. Algorithm

To make the notation consistent across iterations, we introduce a variable α_i for each i in $\{1, \dots, \ell\}$ and restate $\text{LP}_{\text{mcsndp}}(G, f, \{c^i\}_{i=1}^\ell)$ as $\text{LP}_{\text{mcsndp}}(G, f, \{c^i\}_{i=1}^\ell, \{\alpha_i\}_{i=1}^\ell)$:

$$\begin{aligned} & \text{minimize } z \\ & \text{subject to } \sum_{e \in \delta_G(S)} x_e \geq f(S) \quad \forall S \subseteq V \\ & \quad \sum_{e \in E(G)} c_e^i x_e \leq z - \alpha_i \quad \forall i \in \{1, \dots, \ell\} \\ & \quad z \geq 0 \\ & \quad 1 \geq x_e \geq 0 \quad \forall e \in E(G). \end{aligned}$$

Algorithm 1 Iterative Multi-Criteria Network Design Algorithm.

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 $F = \emptyset$, $\alpha_i = 0$ for all $i \in [\ell]$.
2. While $E(G) \neq \emptyset$,
 - (a) Find an optimal extreme point solution, x to $\text{LP}_{\text{mcsndp}}(G, f, \{c^i\}_{i=1}^\ell, \{\alpha_i\}_{i=1}^\ell)$.
 - (b) Delete every edge e with $x_e = 0$ from $E(G)$.
 - (c) Let $E_{1/\ell^+} = \{e \in E(G) : x_e \geq 1/\ell\}$.
 - For every $i \in \{1, \dots, \ell\}$, update $\alpha_i \leftarrow \alpha_i + \sum_{e \in E_{1/\ell^+}} c_e^i x_e$.
 - For every $S \subseteq V(G)$, update $f(S) \leftarrow f(S) - |\delta_G(S) \cap E_{1/\ell^+}|$.
 - Set $F \leftarrow F \cup E_{1/\ell^+}$ and delete E_{1/ℓ^+} from $E(G)$.
3. Return F .

The next lemma implies Theorem 3. It crucially uses Theorem 5 about the existence of a large edge value in an extreme point solution.

Lemma 11. *Let (z^*, x^*) be an optimal extreme point of $\text{LP}_{\text{mcsndp}}(G, f, \{c^i\}_{i=1}^\ell, \{\alpha_i\}_{i=1}^\ell)$ and F be a solution returned by Algorithm 1 on input $(G, f, \{c^i\}_{i=1}^\ell)$, then*

$$\sum_{e \in F} c_e^i \leq \ell \cdot (z^* - \alpha_i) \quad \forall i \in \{1, \dots, \ell\}.$$

Proof. For any extreme point solution of $\text{LP}_{\text{mcsndp}}(G, f, \{c^i\}_{i=1}^\ell, \{\alpha_i\}_{i=1}^\ell)$, $E_{1/\ell^+} \neq \emptyset$ from Theorem 5. So, the size of the set of edges, $E(G)$, is strictly decreasing in every iteration of Algorithm 1. We will prove the lemma by induction. After the first iteration of Algorithm 1, let E_{res} denote the set of edges that are not deleted, then the residual LP is:

$$\begin{aligned}
& \text{minimize } z \\
& \text{subject to } \sum_{e \in \delta_G(S) \cap E_{res}} x_e \geq f(S) - \sum_{e \in E_{1/\ell^+} \cap \delta_G(S)} 1 \quad \forall S \subseteq V \\
& \quad \sum_{e \in E_{res}} c_e^i x_e \leq z - \alpha_i - \sum_{e \in E_{1/\ell^+}} c_e^i x_e \quad \forall i \in \{1, \dots, \ell\} \\
& \quad z \geq 0 \\
& \quad 1 \geq x_e \geq 0 \quad \forall e \in E_{res}.
\end{aligned}$$

By Lemma 8, the requirement function for the residual LP is skew supermodular. Let z_{res}^* denote the optimal value of the residual LP. By the induction hypothesis, $\sum_{e \in F \setminus E_{1/\ell^+}} c_e^i \leq \ell \cdot (z_{res}^* - \alpha_i - \sum_{e \in E_{1/\ell^+}} c_e^i x_e^*)$. Clearly, z^* and x^* restricted to E_{res} is a feasible solution for the residual LP. Thus, $z^* \leq z_{res}^*$, and

$$\begin{aligned}
\sum_{e \in F} c_e^i &= \sum_{e \in E_{1/\ell^+}} c_e^i + \sum_{e \in F \setminus E_{1/\ell^+}} c_e^i \\
&\leq \sum_{e \in E_{1/\ell^+}} c_e^i (\ell x_e^*) + \ell \cdot \left(z_{res}^* - \alpha_i - \sum_{e \in E_{1/\ell^+}} c_e^i x_e^* \right) \\
&\leq \ell \cdot (z^* - \alpha_i). \quad \square
\end{aligned}$$

Algorithm 1 can be implemented in strongly polynomial time for the generalized Steiner network case using the Algorithm from Section 9 in [11] to solve $LP_{mcsndp}(G, f, \{c^i\}_{i=1}^\ell, \{\alpha_i\}_{i=1}^\ell)$ in strongly polynomial time.

4. Proof of Theorem 4

Let x be an extreme point solution to $LP_{mcsndp}(G, f, \{c^i\}_{i=1}^\ell, \{\alpha_i\}_{i=1}^\ell)$. If $x_e = 0$ for any edge e , then we can delete it and x is still an extreme point solution for the residual LP. So, we can assume $x_e > 0$ for each $e \in E$. Let $\mathcal{B}$ be the laminar family of tight sets from Theorem 10. We form a directed forest, $\mathcal{F}$, whose node set is the laminar family $\mathcal{B}$. $\mathcal{F}$ contains an edge from U to V if U is the smallest set in $\mathcal{F}$ containing V . An edge crosses a node if one of its endpoints is inside the node and the other one is outside the node.

From Theorem 10, we get $|E| - (\ell - 1) \leq |\mathcal{F}| \leq |E| + 1$. We will use the fractional token assignment procedure described in [15]. Each edge $e = (u, v)$ is given 1 token which it distributes as follows:

1. Let S be the smallest set in $\mathcal{F}$ containing u , and R be the smallest set in $\mathcal{F}$ containing v . Then e assigns x_e tokens to each of S and R .
2. Let T be the smallest set in $\mathcal{F}$ containing both u and v . Then e assigns $1 - 2x_e$ tokens to T .

We will prove that if $0 < x_e < 1/\ell$ for every edge e in E , then $|\mathcal{F}| < |E| - (\ell - 1)$, a contradiction.

Lemma 12. [15] Every node in $\mathcal{F}$ contains at least 1 token.

For the reader's convenience, we restate the proof from [15].

Proof. For a node $U \in \mathcal{F}$ with children $\{V_i\}_{i=1}^k$, let

$$\begin{aligned}
A &= \left\{ e : |e \cap U| = 1, |e \cap \left(\bigcup_{i=1}^k V_i \right)| = 0 \right\}, \\
B &= \left\{ e : |e \cap U| = 2, |e \cap \left(\bigcup_{i=1}^k V_i \right)| = 1 \right\},
\end{aligned}$$

$$C = \left\{ e : |e \cap U| = 2, |e \cap \left(\bigcup_{i=1}^k V_i \right)| = 2 \right\}.$$

Since $U, \{V_i\}_{i=1}^k$ are tight sets,

$$\begin{aligned}
f(U) - \sum_{i=1}^k f(V_i) &= x(\delta(U)) - \sum_{i=1}^k x(\delta(V_i)) \\
&= x(A) - x(B) - 2x(C).
\end{aligned}$$

Note that $A \cup B \cup C \neq \emptyset$, otherwise $x(\delta(U)) = \sum_{i=1}^k x(\delta(V_i))$. The set U gets x_e tokens from every edge e in A , $1 - x_e$ tokens from every edge e in B and $1 - 2x_e$ tokens from every edge e in C . So, the total amount of tokens received by U equals

$$\begin{aligned}
\sum_{e \in A} x_e + \sum_{e \in B} (1 - x_e) + \sum_{e \in C} (1 - 2x_e) \\
&= |B| + |C| + x(A) - x(B) - 2x(C) \\
&= |B| + |C| + f(U) - \sum_{i=1}^k f(V_i)
\end{aligned}$$

which is positive due to the LHS and integral due to the RHS and hence at least one. $\square$

We proceed to the proof of Theorem 5.

Proof. From Lemma 12, every set in $\mathcal{F}$ contains at least 1 token. Now, we will prove that the amount of tokens assigned to nodes in $\mathcal{F}$ is strictly less than $|E| - (\ell - 1)$, which implies that $|\mathcal{F}| < |E| - (\ell - 1)$, a contradiction.

Consider a root node R in $\mathcal{F}$. Every edge $e \in \delta(R)$ does not assign the token corresponding to step 2 to any set in $\mathcal{F}$. The total amount of unassigned tokens is

$$\sum_{e \in \delta(R)} (1 - 2x_e) = |\delta(R)| - 2f(R).$$

This quantity is integral and strictly greater than 0 because $1 - 2x_e > 0$ and $\delta(R) \neq \emptyset$. Since $f(R) \geq 1$ and $0 < x_e < 1/\ell$ for all e , the number of edges in $\delta(R)$ is at least $\ell + 1$ and $1 - 2x_e > 1 - \frac{2}{\ell}$. So,

$$\begin{aligned}
\sum_{e \in \delta(R)} (1 - 2x_e) &> (\ell + 1) \cdot \left(1 - \frac{2}{\ell} \right) \\
&= \ell + 1 - 2 - \frac{2}{\ell} \\
&= \ell - 1 - \frac{2}{\ell}.
\end{aligned}$$

So, $\sum_{e \in \delta(R)} (1 - 2x_e)$ is an integer strictly greater than $\ell - 1 - \frac{2}{\ell}$. For $\ell \geq 2$,

$$\sum_{e \in \delta(R)} (1 - 2x_e) \geq \ell - 1.$$

Thus $\mathcal{F}$ contains at most $|E| - (\ell - 1)$ tokens. Now if $\delta(R)$ contains an edge e whose other endpoint is not present in $\mathcal{F}$, then $\mathcal{F}$ does not contain the token corresponding to the other end of e , and hence the amount of tokens in $\mathcal{F}$ is at most

$$|E| - (\ell - 1) - x_e < |E| - (\ell - 1).$$

Otherwise, there must be at least one edge e , that does not cross R but crosses another root node in $\mathcal{F}$, else

$$\chi(\delta(R)) = \sum_{V \neq R, V \text{ is a root node}} \chi(\delta(V)),$$

which violates the linear independence of characteristic vectors of the nodes of $\mathcal{F}$. So the amount of tokens in $\mathcal{F}$ is at most

$$|E| - (\ell - 1) - (1 - 2x_e) < |E| - (\ell - 1). \quad \square$$

Acknowledgements

We thank R. Ravi for helpful discussions and Mehrdad Ghadiri for bringing reference [12] to our attention.

Funding

This work was supported in part by the NSF awards AF:2106444, AF:1910423, CCF-2007443, and CCF-2106644.

References

- [1] V. Barichard, X. Gandibleux, V. T'Kindt, *Multiobjective Programming and Goal Programming: Theoretical Results and Practical Applications*, vol. 618, Springer Science & Business Media, 2008.
- [2] A. Bhargale, S. Kopparty, S. Sachdeva, Simultaneous approximation of constraint satisfaction problems, in: *International Colloquium on Automata, Languages, and Programming*, Springer, 2015, pp. 193–205.
- [3] A. Bhargale, S. Khot, S. Kopparty, S. Sachdeva, D. Thiruvengadachari, Near-optimal approximation algorithm for simultaneous max-cut, in: *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, 2018, pp. 1407–1425.
- [4] A. Chinchuluun, P.M. Pardalos, A. Migdalas, L. Pitsoulis, *Pareto Optimality, Game Theory and Equilibria*, Springer, 2008.
- [5] Z. Drezner, S.Y. Nof, On optimizing bin picking and insertion plans for assembly robots, *IIE Trans.* 16 (3) (1984) 262–270.
- [6] S. French, R. Hartley, L. Thomas, D.J. White, *Multi-Objective Decision Making: Based on the Proceedings of a Conference on Multi-Objective Decision Making*, Academic Press, 1983.
- [7] M.X. Goemans, A.V. Goldberg, S. Plotkin, D.B. Shmoys, E. Tardos, D.P. Williamson, Improved approximation algorithms for network design problems, *Tech. rep.*, Cornell University Operations Research and Industrial Engineering, 1995.
- [8] F. Grandoni, R. Ravi, M. Singh, Iterative rounding for multi-objective optimization problems, in: *European Symposium on Algorithms*, Springer, 2009, pp. 95–106.
- [9] F. Grandoni, R. Ravi, M. Singh, R. Zenklusen, New approaches to multi-objective optimization, *Math. Program.* 146 (1) (2014) 525–554.
- [10] H.W. Hamacher, G. Ruhe, Algorithms for max-linear and multi-criteria spanning tree problems, in: *Operations Research'92*, Springer, 1993, pp. 124–127.
- [11] K. Jain, A factor 2 approximation algorithm for the generalized Steiner network problem, *Combinatorica* 21 (1) (2001) 39–60.
- [12] A. Kasperski, P. Zieliński, On the approximability of minmax (regret) network optimization problems, *Inf. Process. Lett.* 109 (5) (2009) 262–266.
- [13] L.C. Lau, R. Ravi, M. Singh, *Iterative Methods in Combinatorial Optimization*, vol. 46, Cambridge University Press, 2011.
- [14] U. Lebrecht, *Max-lineare Zuordnungsprobleme in der Roboteroptimierung*, Master's thesis, Universität Kaiserslautern, Department of Mathematics, 1991.
- [15] V. Nagarajan, R. Ravi, M. Singh, Simpler analysis of LP extreme points for traveling salesman and survivable network design problems, *Oper. Res. Lett.* 38 (3) (2010) 156–160.
- [16] C.H. Papadimitriou, M. Yannakakis, On the approximability of trade-offs and optimal access of web sources, in: *Proceedings 41st Annual Symposium on Foundations of Computer Science*, IEEE, 2000, pp. 86–92.
- [17] R. Ravi, M.V. Marathe, S.S. Ravi, D.J. Rosenkrantz, H.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*, 1993, pp. 438–447.
- [18] S. Samadi, U. Tantipongpipat, J. Morgenstern, M. Singh, S. Vempala, The price of fair PCA: one extra dimension, *arXiv preprint*, arXiv:1811.00103, 2018.
- [19] M. Singh, L.C. Lau, Approximating minimum bounded degree spanning trees to within one of optimal, *J. ACM* 62 (1) (2015) 1–19.
- [20] U. Tantipongpipat, S. Samadi, M. Singh, J.H. Morgenstern, S.S. Vempala, Multi-criteria dimensionality reduction with applications to fairness, *Adv. Neural Inf. Process. Syst.* 32 (2019).