Group-Strategyproof Cost Sharing for Metric Fault Tolerant Facility Location *

Yvonne Bleischwitz[1] and Florian Schoppmann[1,2]

[1] Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany.
{yvonneb,fschopp}@uni-paderborn.de
[2] International Graduate School of Dynamic Intelligent Systems

Abstract. In the context of *general demand* cost sharing, we present the first group-strategyproof mechanisms for the *metric fault tolerant uncapacitated facility location problem*. They are $(3L)$-budget-balanced and $(3L \cdot (1 + \mathcal{H}_n))$-efficient, where L is the maximum service level and n is the number of agents. These mechanisms generalize the seminal *Moulin mechanisms* for *binary demand*. We also apply this approach to the *generalized Steiner problem in networks*.

1 Introduction and Model

Satisfying agents' connectivity requirements at minimum cost is a major challenge in network design problems. In many cases, these problems are NP-hard. The situation gets even more intricate when we consider these problems in the context of *cost sharing*.

In cost-sharing scenarios, a *service provider* offers a common *service* (e.g., connectivity within a network) to agents. Based on *bids* that indicate the agents' willingness to pay, the provider determines a service allocation and payments. Moreover, he computes a solution to provide the service according to the allocation (e.g., to meet the connectivity requirements).

The decision-making of the provider is governed by a commonly known *cost-sharing mechanism*. Essential properties of these mechanisms are *group-strategyproofness*, preventing collusion by guaranteeing that rational agents communicate bids equal to their *true valuations*; *budget-balance*, ensuring recovery of the provider's cost as well as competitive prices in that the generated surplus is always relatively small; and *economic efficiency*, providing a reasonable trade-off between the provider's cost and the valuations of the excluded agents. Finally, practical applications demand for polynomial-time computability (in the size of the problem).

Most research assumes *binary demand*, where agents are "served" or "not served". In contrast, we consider the *general demand* setting, providing *service levels* ranging from 0 to some maximum number. This is of particular interest

* This work was partially supported by the IST Program of the European Union under contract number IST-15964 (AEOLUS).

when agents require different *qualities of service*. For connectivity problems, the service level of an agent is the number of her (distinct) connections. More connections correspond to a higher quality of service, for reasons including throughput and resistance to link failure.

1.1 The Model

Notation. For $n \in \mathbb{N}$, let $[n] := \{1, \dots, n\}$ and $[n]_0 := [n] \cup \{0\}$. Let $\mathcal{H}_n := \sum_{i=1}^{n} \frac{1}{i} \in (\log n, 1 + \log n)$. For $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Q}^n$, we write $\boldsymbol{x} \geq \boldsymbol{y}$, if for all $i \in [n]$, $x_i \geq y_i$. Let $\boldsymbol{0}, \boldsymbol{1}, \boldsymbol{e}_i$ be the zero, one, and i-th standard basis vector. We say that $\boldsymbol{x} \in \{0,1\}^n$ *indicates* $X \subseteq [n]$ if $x_i = 1 \Leftrightarrow i \in X$.

We consider a finite set $[n]$ of agents. Each agent $i \in [n]$ has a *maximum level of service* $L_i \in \mathbb{N}$ she can receive. Let $\mathcal{Q} := [L_1]_0 \times \ldots \times [L_n]_0$, $\mathcal{L} := \mathbb{Q}^{L_1} \times \ldots \times \mathbb{Q}^{L_n}$, $\mathcal{L}_{\geq 0} := \mathbb{Q}^{L_1}_{\geq 0} \times \ldots \times \mathbb{Q}^{L_n}_{\geq 0}$, and $L := \max_{i \in [n]}\{L_i\}$.

An *allocation* is a vector $\boldsymbol{q} \in \mathcal{Q}$ representing the service level given to each agent. Given $\boldsymbol{q} \in \mathcal{Q}$ and $l \in [L]$, we define $\boldsymbol{q}^{\leq l}$ by $q_i^{\leq l} := \min_i\{q_i, l\}$, and let $\boldsymbol{q}^l \in \{0,1\}^n$ indicate the set $Q_l := \{i \in [n] \mid q_i \geq l\}$.

The *valuation vector* $\boldsymbol{v}_i \in \mathbb{Q}^{L_i}$ of agent i consists of the *marginal valuations* $v_{i,l}$ of receiving level l additionally to level $l-1$. Agent i's total valuation for level m is thus $\sum_{l=1}^{m} v_{i,l}$. We call $\boldsymbol{V} = (\boldsymbol{v}_1, \dots, \boldsymbol{v}_n) \in \mathcal{L}$ the *valuation matrix*; accordingly, $\boldsymbol{B} = (\boldsymbol{b}_1, \dots, \boldsymbol{b}_n) \in \mathcal{L}$ denotes a *bid matrix*.

Assumption: We assume that $v_{i,1} \geq \ldots \geq v_{i,L_i}$ for all $i \in [n]$.

Definition 1. *A* cost-sharing mechanism $M = (q, x) : \mathcal{L} \to \mathcal{Q} \times \mathbb{Q}^n_{\geq 0}$ *is a function that gets as input a bid matrix $\boldsymbol{B} \in \mathcal{L}$. It outputs an allocation $q(\boldsymbol{B}) \in \mathcal{Q}$ and a vector of cost shares $x(\boldsymbol{B}) \in \mathbb{Q}^n_{\geq 0}$.*

We always require three standard properties of cost-sharing mechanisms:

- *No positive transfers* (NPT): For all $\boldsymbol{B} \in \mathcal{L}$, it is $x(\boldsymbol{B}) \geq 0$.
- *Voluntary participation* (VP): Agents are never charged more than they bid, i.e., for all $\boldsymbol{B} \in \mathcal{L}$ and all $i \in [n]$, it is $x_i(\boldsymbol{B}) \leq \sum_{l=1}^{q_i(\boldsymbol{B})} b_{i,l}$.
- *Consumer sovereignty* (CS): An agent can always ensure to receive a certain service level, i.e., for all $i \in [n]$ and all $l \in [L_i]_0$, there is a bid vector $\boldsymbol{b}_i^{+l} \in \mathbb{Q}^{L_i}$ such that for all $\boldsymbol{B} \in \mathcal{L}$ with $\boldsymbol{b}_i = \boldsymbol{b}_i^{+l}$, it is $q_i(\boldsymbol{B}) = l$.

An agent i aims to submit a bid vector $\boldsymbol{b}_i$ such that her utility is maximized, where her utility is given as $u_i(\boldsymbol{B}) := \sum_{l=1}^{q_i(\boldsymbol{B})} v_{i,l} - x_i(\boldsymbol{B})$.

Definition 2. *A cost-sharing mechanism $M = (q, x)$ is* group-strategyproof *(GSP) if for every true valuation matrix $\boldsymbol{V} \in \mathcal{L}$ and any coalition $K \subseteq [n]$ there is no bid matrix $\boldsymbol{B} \in \mathcal{L}$ with $\boldsymbol{b}_i = \boldsymbol{v}_i$ for all $i \notin K$, such that $u_i(\boldsymbol{B}) \geq u_i(\boldsymbol{V})$ for all $i \in K$ with at least one strict inequality.*

Cost shares computed by a GSP mechanism only depend on the computed allocation [Mou99]. This gives rise to the following definition:

Definition 3. *A* cost-sharing method $\xi : \mathcal{Q} \to \mathbb{Q}^n_{\geq 0}$ *maps each allocation* $\boldsymbol{q} \in \mathcal{Q}$ *to a vector of cost shares. If* $q_i = 0$*, we require* $\xi_i(\boldsymbol{q}) = 0$*.*

For *binary demand*, *cross-monotonic* cost-sharing methods $\xi : \{0,1\}^n \to \mathbb{Q}^n_{\geq 0}$ are of particular interest for achieving GSP. For all allocations $\boldsymbol{p}$ and all agents $i \in [n]$ and $j \in [n] \setminus \{i\}$ with $p_j = 0$, they fulfill $\xi_i(\boldsymbol{p}) \geq \xi_i(\boldsymbol{p} + \boldsymbol{e}_j)$. Using such a ξ, the (binary demand) mechanism Moulin_ξ is GSP [Mou99]:

Algorithm 1 (computing $\mathit{Moulin}_\xi = (q, x)$**).**
Input: $\boldsymbol{B} \in \mathbb{Q}^n$ ▷ $\boldsymbol{B} = (b_{1,1}, \dots, b_{n,1})$
Output: assignment $q(\boldsymbol{B}) \in \{0,1\}^n$, cost-share vector $x(\boldsymbol{B}) \in \mathbb{Q}^n_{\geq 0}$

1: $\boldsymbol{p} := \boldsymbol{1}$;
2: **while** $\boldsymbol{p} \neq \boldsymbol{0}$ and there exists i with $b_{i,1} < \xi_i(\boldsymbol{p})$ **do**
3: $p_j := 0$ for an arbitrary j with $b_{j,1} < \xi_j(\boldsymbol{p})$
4: **return** $(\boldsymbol{p}, \xi(\boldsymbol{p}))$

Our general demand mechanisms use *marginal* cost-sharing methods:

Definition 4. *A* marginal cost-sharing method *is a function* $\chi : \mathcal{Q} \to \mathcal{L}_{\geq 0}$*, where* $\chi_{i,l}(\boldsymbol{q})$ *is the marginal cost-share of agent* i *for additionally receiving service level* l *to level* $l-1$*. If* $l > q_i$*, we require* $\chi_{i,l}(\boldsymbol{q}) = 0$*.*

We now turn to the service cost. A *cost-sharing problem* is specified by a *cost function* $C : \mathcal{Q} \to \mathbb{Q}_{\geq 0}$ mapping each $\boldsymbol{q}$ to the cost of providing each agent i with service level q_i. Typically, costs stem from solutions to a combinatorial optimization problem and are defined only implicitly. We let $C(\boldsymbol{q})$ be the value of a minimum-cost solution for the instance induced by $\boldsymbol{q}$. This cost can in general not be recovered exactly due to restrictions placed by the GSP requirement. Further difficulties arise when problems are hard. We denote the cost of an approximate solution by $C'(\boldsymbol{q})$ and require the total charge of a mechanism to lie within reasonable bounds:

Definition 5. *A general demand cost-sharing mechanism* $M = (q, x)$ *is* β-budget-balanced *(*β*-BB, for* $\beta \geq 1$*) with respect to* $C, C' : \mathcal{Q} \to \mathbb{Q}_{\geq 0}$ *if for all* $\boldsymbol{B} \in \mathcal{L}$ *it holds that*

$$C'(q(\boldsymbol{B})) \leq \sum_{i=1}^{n} x_i(\boldsymbol{B}) \leq \beta \cdot C(q(\boldsymbol{B})) \ .$$

As a quality measure for the computed allocation, we use *optimal social costs* $SC_{\boldsymbol{V}}(\boldsymbol{q}) := C(\boldsymbol{q}) + \sum_{i=1}^{n} \sum_{l=q_i+1}^{L_i} \max\{0, v_{i,l}\}$ and *actual social costs* $SC'_{\boldsymbol{V}}(\boldsymbol{q}) := C'(\boldsymbol{q}) + \sum_{i=1}^{n} \sum_{l=q_i+1}^{L_i} \max\{0, v_{i,l}\}$. The cost incurred and the valuations of the rejected agents should be traded off as good as possible:

Definition 6. *A general demand cost-sharing mechanism* $M = (q, x)$ *is* γ-efficient *(*γ*-EFF, for* $\gamma \geq 1$*) with respect to* $C, C' : \mathcal{Q} \to \mathbb{Q}_{\geq 0}$ *if for all true valuations* $\boldsymbol{V} \in \mathcal{L}$ *it holds that* $SC'_{\boldsymbol{V}}(q(\boldsymbol{V})) \leq \gamma \cdot \min_{\boldsymbol{q} \in \mathcal{Q}}\{SC_{\boldsymbol{V}}(\boldsymbol{q})\}$*.*

The efficiency of Moulin_ξ can be analyzed via the *summability* of ξ:

Definition 7. *A binary demand cost-sharing method* $\xi : \{0,1\}^n \to \mathbb{Q}^n_{\geq 0}$ *is* α*-summable* (α*-SUM, for* $\alpha \geq 1$) *with respect to* $C : \mathcal{Q} \to \mathbb{Q}_{\geq 0}$ *if for every* $\boldsymbol{s} \in \{0,1\}^n$ *and every ordering* $s_1, \dots, s_{|S|}$ *of* $S := \{i \in [n] \mid s_i = 1\}$*, it is* $\sum_{i=1}^{|S|} \xi_{s_i}(\boldsymbol{s}^i) \leq \alpha \cdot C(\boldsymbol{s})$*, where* $\boldsymbol{s}^i \in \{0,1\}$ *indicates* $S_i := \{s_1, \dots, s_i\}$*.*

If ξ is β-BB and α-SUM, then $Moulin_\xi$ is $(\alpha+\beta)$-EFF [RS06].

1.2 The Problems

FaultTolerantFL (Metric Fault Tolerant Uncapacitated Facility Location Problem): An instance of this problem is given by a set of agents $[n]$, a set F of facilities, an opening cost $o_f \in \mathbb{N}$ for each $f \in F$, and a non-negative cost function $c : ([n] \cup F) \times ([n] \cup F) \to \mathbb{N}$ that satisfies the triangle inequality. Given $\boldsymbol{q} \in \mathcal{Q}$ with $\max_i\{q_i\} \leq |F|$, the aim is to open a set of facilities and connect each agent i to q_i distinct open facilities, such that the total opening and connection cost is minimized. For $k \in \mathbb{N}$, let $\mathcal{F}_k := \{F' \subseteq F \mid |F'| \geq k\}$. For $F' \in \mathcal{F}_{\max_i\{q_i\}}$, let F'_i be a set of q_i closest facilities in F' to $i \in [n]$. Then the optimal cost is $C(\boldsymbol{q}) := \min_{F' \in \mathcal{F}_{\max_i\{q_i\}}}\{\sum_{f \in F'} o_f + \sum_{i \in [n]} \sum_{f \in F'_i} c(i,f)\}$.

GeneralizedSteiner (Generalized Steiner Problem in Networks): An instance of this problem is given by a set of agents $[n]$, an undirected graph $G := (V,E)$ with edge costs $c : E \to \mathbb{N}$, and a pair of nodes (s_i, t_i) for each agent $i \in [n]$. For a requirement vector $\boldsymbol{q}$, the aim is to determine a minimum-cost subgraph with cost $C(\boldsymbol{q})$ that has q_i edge-disjoint paths between s_i and t_i. We consider a simplification allowing to use multiple edge copies. The cost of such an edge copy is equal to the cost of the edge.

1.3 Related Work

The only general design technique for binary demand GSP mechanisms is applying $Moulin_\xi$ [Mou99, JV01]. Particularly, cross-monotonic cost-sharing methods were designed by Pál et al. [PT03] for metric uncapacitated facility location and by Könemann et al. [KLS05] for Steiner forests. A non-Moulin and only SP mechanism for facility location was introduced by Devanur et al. [DMV05].

Prior to this work, incremental cost-sharing mechanisms were the only known GSP mechanisms for general demand cost sharing. They simply consider an ordering that specifies which agent's level is incremented next and make the agent pay for the corresponding marginal cost. Incremental mechanisms are only known to be GSP for certain costs, and are essentially the only GSP and 1-BB mechanisms for these costs [Mou99]. However, for costs induced by FaultTolerantFL and GeneralizedSteiner, incremental mechanisms are not GSP. The interested reader can find examples in the extended version of this paper.

The only other general demand mechanisms we are aware of are *acyclic mechanisms* introduced by Mehta et al. [MRS07], which were originally designed for binary demand to overcome the limitations of Moulin mechanisms [IMM05, RS06]. The main drawback of acyclic mechanisms is that they are only *weak* GSP, meaning that coalitions are only assumed to be successful if *every* agent strictly improves her utility. Technically, *all* inequalities in Definition 2 are strict. For FaultTolerantFL, [MRS07] give a $O(L^2)$-BB and $O(L^2 \cdot (1+\log n))$-EFF acyclic mechanism. In the full version of their paper, they present a $\mathcal{H}_n$-BB and $(2\mathcal{H}_n \cdot (1 + \mathcal{H}_L))$-EFF acyclic mechanism for the non-metric case.[3]

To the best of our knowledge, the best approximation algorithm for FaultTolerantFL yields an appproximation factor of 2.076 [SSL03], while the best approximation factor for GeneralizedSteiner is 2, with and without the simplification of allowing edge copies [Jai01].

1.4 Contribution

The point of departure for this work is a rather obvious idea for generalizing Moulin mechanisms: Start with the maximum allocation and iteratively reduce service levels until every agent can afford her remaining levels. The cost shares are extracted from marginal cost-sharing methods χ. These mechanisms, termed $\mathit{MoulinGD}_\chi$, are stated in Section 2.

- We identify three properties of marginal cost-sharing methods that are sufficient for $\mathit{MoulinGD}_\chi$ to be GSP. It comes as no surprise that a generalization of binary-demand cross-monotonicity is among them.
- We give marginal cost-sharing methods χ^{FL} for every instance of FaultTolerantFL in Section 3 and show that $\mathit{MoulinGD}_{\chi^{FL}}$ is GSP, $(3L)$-BB and $(3L \cdot (1+\mathcal{H}_n))$-EFF. These are the first GSP mechanisms for this problem. Method χ^{FL} is a natural generalization of the binary demand cost-sharing method for facility location in [PT03]. In contrast, the generalization used within acyclic mechanisms in [MRS07] does not guarantee GSP.
- We give the first GSP mechanisms for GeneralizedSteiner in Section 4.

Our work adapts the common assumption that marginal valuations are non-increasing in the service level. Omitted proofs are given in the extended version.

2 Generalized Moulin Mechanisms

Given a marginal cost-sharing method χ, we propose to generalize Moulin mechanisms as in Algorithm 2:

Algorithm 2 (computing $\mathit{MoulinGD}_\chi := (q, x)$).
Input: bid matrix $\boldsymbol{B} \in \mathcal{L}$
Output: allocation $q(\boldsymbol{B}) \in \mathcal{Q}$; cost-share vector $x(\boldsymbol{B}) \in \mathbb{Q}^n_{\geq 0}$

```
q := (L_1, ..., L_n);
while there exists i with b_{i,q_i} < χ_{i,q_i}(q) do
    q_j := q_j − 1 for an arbitrary j with b_{j,q_j} < χ_{j,q_j}(q)
return (q, x) with x_i := Σ_{l=1}^{q_i} χ_{i,l}(q)
```

[3] Note that these results are adjusted to our notion of β-BB.

We state three properties of χ that are sufficient for $MoulinGD_\chi$ to be GSP. The first is a generalization of binary demand cross-montonicity:

Definition 8. *A marginal cost-sharing method is* cross-monotonic *if for all allocations $\boldsymbol{q} \in \mathcal{Q}$, all agents $i \in [n]$ and $j \in [n] \setminus \{i\}$ with $q_j < L_j$, and all service levels $l \in [L_i]$, it holds that $\chi_{i,l}(\boldsymbol{q}) \geq \chi_{i,l}(\boldsymbol{q} + \boldsymbol{e}_j)$.*

The second property ensures that the marginal cost-share $\chi_{i,l}(\boldsymbol{q})$ of agent i with $q_i \geq l$ is exactly the marginal cost-share $\chi_{i,l}(\boldsymbol{q}^{\leq l})$.

Definition 9. *A marginal cost-sharing method is* level-restricted *if for all allocations $\boldsymbol{q} \in \mathcal{Q}$, for all agents $i \in [n]$, and for all service levels $l \in [L_i]$, it holds that $\chi_{i,l}(\boldsymbol{q}) = \chi_{i,l}(\boldsymbol{q}^{\leq l})$.*

The third property together with cross-monotonicity implies that the marginal cost-share of an agent is non-decreasing in the number of levels:

Definition 10. *A marginal cost-sharing method is* non-decreasing *if for all allocations $\boldsymbol{q} \in \mathcal{Q}$, for service level $l := \max_i\{q_i\}$, and for all agents $i \in [n]$ with $q_i = l < L_i$, it holds that $\chi_{i,l}(\boldsymbol{q}) \leq \chi_{i,l+1}(\boldsymbol{q} + \sum_{j \in [n]: q_j = l < L_j} \boldsymbol{e}_j)$.*

Lemma 1. *If χ is non-decreasing and cross-monotonic, it holds for all allocations $\boldsymbol{q} \in \mathcal{Q}$, for all service levels $l \in [L]$, and for all agents $i \in [n]$ with $q_i > l$ that $\chi_{i,l}(\boldsymbol{q}^{\leq l}) \leq \chi_{i,l+1}(\boldsymbol{q}^{\leq l+1})$.*

If χ is level restricted, agent i's utility is the sum over the *marginal utilities* $v_{i,l} - \chi_{i,l}(\boldsymbol{q}^{\leq l})$. If χ is additionally cross-monotonic and non-decreasing, by Lemma 1 and non-increasing marginal valuations, these marginal utilities are non-increasing in l. The proof of Theorem 1 heavily relies on non-increasing marginal utilities and uses the main idea from [MS01], which shows that $Moulin_\xi$ is GSP if ξ is cross-monotonic.

Theorem 1. *$MoulinGD_\chi$ is GSP given any level-restricted, cross-monotonic, and non-decreasing marginal cost-sharing method χ. Furthermore, it satisfies NPT, VP and CS.*

Due to space limitations we only show that if exactly one of the three properties required for χ does not hold, $MoulinGD_\chi$ is not GSP anymore. It remains an open problem whether there are cost functions for which all three properties need to be fulfilled at once in order to obtain GSP.

For all examples, let $n = 2$, $L_1 = L_2 = 2$, and $\chi_{2,l}(\boldsymbol{q}) := 2$ for all $l \in [2]$ and all $\boldsymbol{q} \in \mathcal{Q}$ with $q_2 \geq l$. We always assume that $\boldsymbol{v}_2 = (2, 2)$.

Example 1. Consider χ with $\chi_{1,l}(\boldsymbol{q}) := 1$ for all $l \in [2]$ and all $\boldsymbol{q} \in \mathcal{Q}$ with $q_1 \geq l$, with the only exception that $\chi_{1,2}(2,2) := 2$. Obviously, χ is level-restricted and non-decreasing, but not cross-monotonic since $\chi_{1,2}(2,1) < \chi_{1,2}(2,2)$. For the case that $\boldsymbol{v}_1 = (2,2)$, both agents get service level 2, where $u_1(\boldsymbol{v}_1, \boldsymbol{v}_2) = 1$ and $u_2(\boldsymbol{v}_1, \boldsymbol{v}_2) = 0$. Agent 2 may then bid $\boldsymbol{b}_2 = (-1,-1)$ in order to not receive the service with the result that agent 1 receives level 2 with utility $u_1(\boldsymbol{v}_1, \boldsymbol{b}_2) = 2$.

Example 2. Consider χ with $\chi_{1,1}(\boldsymbol{q}) := 1$ for all $\boldsymbol{q} \in \mathcal{Q}$ with $q_1 = 1$, $\chi_{1,1}(\boldsymbol{q}) := 2$ for all $\boldsymbol{q} \in \mathcal{Q}$ with $q_1 = 2$ and $\chi_{1,2}(\boldsymbol{q}) := 3$ for all $\boldsymbol{q} \in \mathcal{Q}$ with $q_1 = 2$. Here, χ is cross-monotonic and non-decreasing but fails to be level-restricted due to $\chi_{1,1}(2,2) > \chi_{1,1}(1,1)$. If now $\boldsymbol{v}_1 = (3,3)$, then agent 1 receives level 2 and $u_1(\boldsymbol{v}_1, \boldsymbol{v}_2) = 1$. However, for $\boldsymbol{b}_1 = (2,2)$, agent 1 receives only level 1, and $u_1(\boldsymbol{b}_1, \boldsymbol{v}_2) = 2$.

On the other hand, we get the situation $\chi_{1,1}(2,2) < \chi_{1,1}(1,1)$ when we change χ such that $\chi_{1,1}(\boldsymbol{q}) = 2$ for all $\boldsymbol{q} \in \mathcal{Q}$ with $q_1 = 1$ and $\chi_{1,1}(\boldsymbol{q}) = 1$ for all $\boldsymbol{q} \in \mathcal{Q}$ with $q_1 = 2$. If $\boldsymbol{v}_1 = (3, 3-\varepsilon)$, then agent 1 receives only one level and $u_1(\boldsymbol{v}_1, \boldsymbol{v}_2) = 1$. However, she may bid $\boldsymbol{b}_1 = (3,3)$ to receive both levels such that $u_1(\boldsymbol{b}_1, \boldsymbol{v}_2) = 2 - \varepsilon$.

Example 3. Consider χ with $\chi_{1,1}(\boldsymbol{q}) := 2$ for all $\boldsymbol{q} \in \mathcal{Q}$ with $q_1 \geq 1$ and $\chi_{1,2}(\boldsymbol{q}) := 1$ for all $\boldsymbol{q} \in \mathcal{Q}$ with $q_1 \geq 2$. Now we have the case that χ is cross-monotonic and level-restricted, but not non-decreasing. For $\boldsymbol{v}_1 = (1,1)$, agent 1 receives level 2 and has a utility of $u_1(\boldsymbol{v}_1, \boldsymbol{v}_2) = -1$. However, bidding $\boldsymbol{b}_1 = (-1,-1)$ ensures a utility of zero.

3 Metric Fault Tolerant Uncapacitated Facility Location

We explain how to define $\chi^{FL} : \mathcal{Q} \to \mathcal{L}_{\geq 0}$ and how to construct a solution in polynomial time. For $\boldsymbol{q} \in \{0,1\}^n$, both reduces to the method and solution by Pál and Tardos [PT03] for binary demand facility location.

Fix $\boldsymbol{q} \in \mathcal{Q}$ and $l \in [L]$. We only need to determine $\chi^{FL}_{i,l}(\boldsymbol{q})$ for all $i \in Q_l := \{j \in [n] \mid q_j \geq l\}$. Simultaneously, every agent i in Q_l uniformly grows a ball with i at its center to infinity. This ball, the *ghost* of i, has radius t at time t. We say that the ghost of i *touches* facility f at time t if $c(i,f) \leq t$. After the ghost of i touches f it starts *filling* f, contributing $t - c(i,f)$ at time $t \geq c(i,f)$. Facility f is said to be *full* if all such contributions sum up to its opening cost o_f. Let $t(f)$ denote the time when f becomes full, and $S_f := \{i \in [n] \mid c(i,f) < t(f)\}$ the set of agents that contributed to filling f. It holds that $\sum_{i \in S_f}(t(f) - c(i,f)) = o_f$.

We define $\chi^{FL}_{i,l}(\boldsymbol{q})$ to be the time that it takes for the ghost of i to touch l full facilities. Note that $\chi^{FL}_{*,l}(\boldsymbol{q})$ only depends on $Q_l := \{j \in [n] \mid q_j \geq l\}$. For $\boldsymbol{q}^l \in \{0,1\}^n$ indicating Q_l, it is $\chi^{FL}_{i,l}(\boldsymbol{q}) := \chi^{FL}_{i,l}(l \cdot \boldsymbol{q}^l)$. This is even stronger than level-restriction.

Lemma 2. *χ^{FL} is level-restricted, cross-monotonic, and non-decreasing.*

We construct a solution with cost $C'(\boldsymbol{q})$ iteratively by computing $\chi^{FL}_{i,l}(\boldsymbol{q})$ for all $i \in Q_l$ for $l = 1, \ldots, \max_i\{q_i\}$. Fix an iteration l. Let $t(f)$ and S_f be the values obtained for all $f \in F$ by growing the ghosts of Q_l.

Facilities are opened in iteration l according to the following rule: Let F_{l-1} be the set of the already opened facilities in iterations $1, \ldots, l-1$. If in iteration l, a facility $f \notin F_{l-1}$ becomes full, we open f if and only if conditions O_1 and O_2 hold:

O_1) There is no facility g that was already opened in iteration l and for which $c(g, f) \leq 2 \cdot t(f)$.

O_2) There are no l distinct facilities $g_1, \dots, g_l \in F_{l-1}$ for which $c(g_k, f) \leq 2 \cdot t(f)$ for all $k \in [l]$.

In order to simplify the analysis, we connect every agent $i \in Q_l$ to one (more) open facility in each iteration l, rather than connecting her with the q_i closest facilities in the final facility set. This is how we do it:

C_1) If $i \in S_f$ for an f opened in iteration l, we connect i to f. (It can be shown that for f, g with $f \neq g$ opened in iteration l, $S_f \cap S_g = \emptyset$.)

C_2) Otherwise, if at time $\chi^{FL}_{i,l}(\boldsymbol{q})$ the ghost of i touches an arbitrary open facility f to which i is not connected yet, we connect i to f.

C_3) Otherwise, let f be an arbitrary full but closed facility that the ghost of i touches at time $\chi^{FL}_{i,l}(\boldsymbol{q})$; f was not opened in iteration l, because O_1 or O_2 do not hold:

a) If O_1 does not hold because of facility g, connect i to g.

b) If O_2 does not hold because of facilities $g_1, \dots, g_l$, connect i to a $g \in \{g_1, \dots, g_l\}$ to which i is not connected yet.

If there are ties in C_3a and C_3b, or when facilities become full simultaneously, break them arbitrarily. It is straightforward to see that for every instance of FaultTolerantFL, χ^{FL} and the solution constructed above can be computed in polynomial time (in the size of the input). Deleting O_2 and C_3b, we get the rules of [PT03]. However, these two rules are crucial for a reasonable BB approximation. We refer to the extended version for the details.

Example 4. Facilities are illustrated as houses, where the roofs are labeled with the opening cost. Agents are circles labeled with the agent's identity. Edges (i, j) are labeled with $c(i, j)$. If i and j are not directly linked, $c(i, j)$ is defined as the cost of a shortest path between i and j in the network.

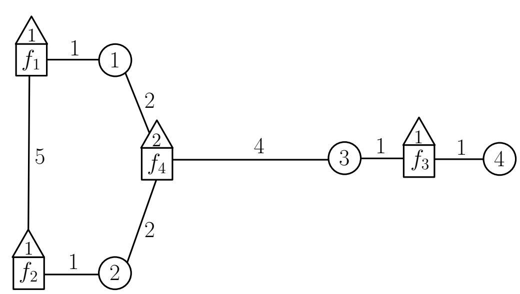

We look at allocation $\boldsymbol{q} = (2, 2, 2, 1)$. The marginal cost shares for level 1 are determined by growing the ghosts of agent set $Q_1 = \{1, 2, 3, 4\}$, where $\boldsymbol{q}^1 = (1, 1, 1, 1)$. For level 2, we grow the ghosts of agent set $Q_2 = \{1, 2, 3\}$, where $\boldsymbol{q}^2 = (1, 1, 1, 0)$. We write $i \circ f$ if i touches f, but f is not full yet; **f** if f becomes full; $i \bullet \mathbf{f}$ if i touches a full facility f. Then for timesteps t:

t	Events for Level 1
1	$1 \circ f_1, 2 \circ f_2, 3 \circ f_3, 4 \circ f_3$
$\frac{3}{2}$	$\mathbf{f_3}, 3 \bullet \mathbf{f_3}, 4 \bullet \mathbf{f_3}$
2	$\mathbf{f_1}, \mathbf{f_2}, 1 \bullet \mathbf{f_1}, 2 \bullet \mathbf{f_2}, 1 \circ f_4, 2 \circ f_4$

t	Events for Level 2
1	$1 \circ f_1, 2 \circ f_2, 3 \circ f_3$
2	$\mathbf{f_1}, \mathbf{f_2}, \mathbf{f_3}, 1 \bullet \mathbf{f_1}, 2 \bullet \mathbf{f_2}, 3 \bullet \mathbf{f_3}$ $1 \circ f_4, 2 \circ f_4$
3	$\mathbf{f_4}, 1 \bullet \mathbf{f_4}, 2 \bullet \mathbf{f_4}$
4	$3 \bullet \mathbf{f_4}$

Note that $i \in S_f$ iff event $i \circ f$ occurs at a strictly smaller time step than event $i \bullet \mathbf{f}$. The cost shares are $\chi_{*,1}(\boldsymbol{q}) = \chi_{*,1}(1,1,1,1) = (2,2,\frac{3}{2},\frac{3}{2})$ and $\chi_{*,2}(\boldsymbol{q}) = \chi_{*,2}(2,2,2,0) = (3,3,4,0)$. The final cost shares are thus $(5,5,\frac{11}{2},\frac{3}{2})$. For Level 1, we open f_3, f_1, and f_2. Due to C_1, we connect agents 3 and 4 to f_3, agent 1 to f_1, and agent 2 to f_2. For Level 2, f_4 stays closed due to O_2. All agents in $\{1,2,3\}$ are connected due to C_3b; 1 is connected to f_2, 2 is connected to f_1, and 3 is connected to f_1 or f_2.

Theorem 2. *For any $\boldsymbol{q} \in \mathcal{Q}$, $X(\boldsymbol{q}) := \sum_{l=1}^{L} \sum_{i \in Q_l} \chi_{i,l}^{FL}(\boldsymbol{q}) \leq L \cdot C(\boldsymbol{q})$. Furthermore, there exists a solution for allocation $\boldsymbol{q}$ with cost $C'(\boldsymbol{q})$, such that $\frac{1}{3} \cdot C'(\boldsymbol{q}) \leq X(\boldsymbol{q})$.*

Proof. Let $\chi := \chi^{FL}$. Fix $\boldsymbol{q} \in \mathcal{Q}$. We show the upper bound; the lower bound is obtained by modifying the proof in [PT03]. Consider an arbitrary facility set $F' \subseteq \mathcal{F}_{\max_i\{q_i\}}$. Fix $l \in [\max_i\{q_i\}]$ and $i \in Q_l$. Let $F'_i \subseteq F'$ be a set of q_i distinct closest facilities in F' to i. We show:

$$\exists f \in F'_i : \chi_{i,l}(\boldsymbol{q}) \leq \begin{cases} t(f) & \text{if } i \in S_f \\ c(i,f) & \text{otherwise .} \end{cases} \qquad (1)$$

Assume that (1) does not hold. Then for all $f \in F'_i$ it holds that $\chi_{i,l}(\boldsymbol{q}) > t(f) > c(i,f)$ if $i \in S_f$ and $\chi_{i,l}(\boldsymbol{q}) > c(i,f) \geq t(f)$ otherwise. Thus, at time $t := \max_{f \in F'_i}\{t(f), c(i,f)\}$, the ghost of i touches at least $q_i \geq l$ full facilities, a contradiction to $t < \chi_{i,l}(\boldsymbol{q})$. Note that (1) especially holds for $F' = F^*$, when F^* is an optimal facility set for $\boldsymbol{q}$. Then,

$$\begin{aligned} \sum_{i \in Q_l} \chi_{i,l}(\boldsymbol{q}) &\leq \sum_{i \in [n]} \left(\sum_{f \in F^*_i : i \in S_f} t(f) + \sum_{f \in F^*_i : i \notin S_f} c(i,f) \right) \\ &= \sum_{f \in F^*} \sum_{i \in S_f : f \in F^*_i} (t(f) - c(i,f)) + \sum_{i \in [n]} \sum_{f \in F^*_i} c(i,f) \\ &\leq \sum_{f \in F^*} o_f + \sum_{i \in [n]} \sum_{f \in F^*_i} c(i,f) = C(\boldsymbol{q}) \ . \end{aligned}$$

Finally, $X(\boldsymbol{q}) = \sum_{l=1}^{L} \sum_{i \in Q_l} \chi_{i,l}(\boldsymbol{q}) \leq L \cdot C(\boldsymbol{q})$. □

Theorem 3. *MoulinGD*$_{\chi^{FL}}$ *is* $(3L \cdot (1 + \mathcal{H}_n))$*-EFF with respect to* C *and* C'.

Whereas we refer to the full version for the proof of Theorem 3, we show a property of χ^{FL} similar to binary demand summability (see e.g. [RS06]) in Lemma 3, which constitutes the main part of the proof.

Lemma 3. *For any* $\boldsymbol{q} \in \mathcal{Q}$ *and any ordering* $s_1, \ldots, s_{|S|}$ *of the set* $S := \{i \in [n] \mid q_i > 0\}$*, where* $\boldsymbol{s}^j \in \{0,1\}^n$ *indicates* $S_j := \{s_1, \ldots, s_j\}$*, it is*

$$\sum_{j=1}^{|S|} \chi^{FL}_{s_j, q_{s_j}}(q_{s_j} \cdot \boldsymbol{s}^j) \leq \mathcal{H}_n \cdot C(\boldsymbol{q}) \ .$$

Proof. Let $\chi := \chi^{FL}$. Roughly speaking, the main idea of the proof is a “reduction” to the summability of a (binary demand) cost-sharing method $\xi^{FL} : \{0,1\}^n \to \mathbb{Q}^n_{\geq 0}$ that we define according to Pál and Tardos [PT03] with a new facility location instance: It has agent set $[n]$, facility set G, and a new cost function $d : ([n] \cup G) \times ([n] \cup G) \to \mathbb{Q}_{\geq 0}$. Let $D := \{0,1\}^n \to \mathbb{Q}_{\geq 0}$ be the optimal cost function for the new instance.

Fix $\boldsymbol{q} \in \mathcal{Q}$ and an ordering $s_1, \ldots, s_{|S|}$ of $S := \{i \in [n] \mid q_i > 0\}$. Fix $j \in [|S|]$ and look at $\chi_{s_j, q_{s_j}}(q_{s_j} \cdot \boldsymbol{s}^j)$, computed for the original instance. For all $f \in F$, let $t(f)$ and S_f be the corresponding values for growing the ghosts of set S_j. In the original instance, let F^* be an optimal facility set for $\boldsymbol{q}$, and $F^*_{s_j}$ be the facilities that s_j is connected to in an optimal solution. It is $F^* \subseteq \mathcal{F}_{\max_i\{q_i\}}$. In the proof of Theorem 2, we have already shown that there exists $f_j \in F^*_{s_j}$, such that $\chi_{s_j, q_{s_j}}(q_{s_j} \cdot \boldsymbol{s}^j)$ is at most $t(f_j)$ if $i \in S_{f_j}$, or $c(s_j, f_j)$ otherwise.

Let the new facilities be $G := \{f_1, \ldots, f_{|S|}\}$. For a each pair in $\{(s_j, f_j)\}_{j \in [|S|]}$, let $d(s_j, f_j) := c(s_j, f_j)$. Furthermore, for all j, j' such that $f_j = f_{j'}$, we define $d(s_j, s_{j'}) := c(s_j, s_{j'})$. All other costs are defined to be sufficiently large, while ensuring that d satisfies the triangle inequality. The networks below illustrate the old (left) and the new (right) facility instance. The grey parts correspond to the unchanged distances.

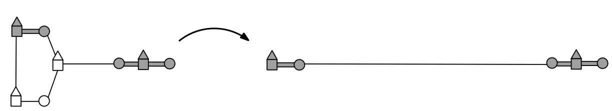

By construction of the new instance, it is $\chi_{s_j, q_{s_j}}(q_{s_j} \cdot \boldsymbol{s}^j) \leq \xi^{FL}_{s_j}(\boldsymbol{s}^j)$ for all $j \in [|S|]$, where $\xi^{FL}_{s_j}(\boldsymbol{s}^j)$ is computed on the new instance. Additionally, $D(\boldsymbol{s}) \leq C(\boldsymbol{q})$, with $\boldsymbol{s} \in \{0,1\}^n$ indicating S. We further use the fact that ξ^{FL} is $\mathcal{H}_n$-SUM [RS07] in order to obtain

$$\sum_{j=1}^{|S|} \chi^{FL}_{s_j, q_{s_j}}(q_{s_j} \cdot \boldsymbol{s}^j) \leq \sum_{j=1}^{|S|} \xi^{FL}_{s_j}(\boldsymbol{s}^j) \leq \mathcal{H}_n \cdot D(\boldsymbol{s}) \leq \mathcal{H}_n \cdot C(\boldsymbol{q}) \ .$$

□

Corollary 1. *There is a marginal cost-sharing method* χ *such that* $MoulinGD_\chi$ *is* $(3L)$*-BB and* $(3L \cdot (1 + \mathcal{H}_n))$*-EFF with respect to* C *and* C'*.*

Proof. Define χ by multiplying χ^{FL} by 3. Adjusting the proofs of Theorem 2 and Theorem 3 leads to the stated BB and EFF guarantees.

We shortly describe the marginal cost-sharing method χ^{AFL} used in the acyclic mechanism introduced by Mehta et al. [MRS07]. The mechanism itself is essentially Mechanism $MoulinGD_\chi$, where line 2 is replaced by "while there exists i with $b_{i,l} < \chi_{i,l}(\boldsymbol{q})$ for an $l \in [q_i]$". The main difference between χ^{AFL} and χ^{FL} is that $\chi^{FL}_{i,l}(\boldsymbol{q})$ is independent of connections computed in iterations 1 to $l-1$.

The marginal cost shares $\chi^{AFL}_{i,l}(\boldsymbol{q})$ for all $i \in Q_l$ are computed iteratively for $l = 1, \ldots, L$. In each iteration l, the cross-monotonic cost-sharing method of Pál and Tardos [PT03] is invoked for all agents with $q_i \geq l$. The instance is changed in such a way that opening costs are set to 0 for already opened facilities. In order to ensure that each agent i is connected to q_i distinct facilities, the distance $c(i, f)$ for already existing connections between i and f is set to infinity. All other distances stay the same. Consider the network below, given by [MRS07]:

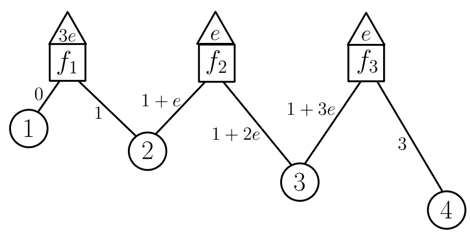

Here, $\chi^{AFL}_{4,2}((2,2,2,2)) = 5 + 5\varepsilon > 3 = \chi^{AFL}_{4,2}(0,2,2,2)$ obviously violates cross-monotonicity. Essentially, this happens due to that fact that for $(2,2,2,2)$, $c(4,3)$ is set to infinity in the first iteration, making her ghost grow longer in the second iteration. However, for $(0,2,2,2)$ it is $c(4,3) = 3$ in both iterations.

4 Generalized Steiner Problem in Networks

For binary demand, Könemann et al. [KLS05] give a polynomial-time computable cross-monotonic cost-sharing method $\xi^{GS} : \{0,1\}^n \to \mathbb{Q}^n_{\geq 0}$ which is $(2 - \frac{1}{n})$-BB. The computed solution is a Steiner forest which can be deduced from the cost-share computation. Additionally, it is known that ξ^{GS} is $O(\log^2 n)$-SUM [CRS06]. The cost-sharing method can essentially be computed by the algorithm by Agrawal et al. [AKR95] with only a small modification which is crucial for cross-monotonicity.

We combine the cost-sharing method from [KLS05] with a straightforward solution construction by Goemans and Bertsimas [GB93]. The marginal cost-sharing method $\chi^{GS} : \mathcal{Q} \to \mathcal{L}_{\geq 0}$ is simply defined by letting $\chi^{GS}_{i,l}(\boldsymbol{q}) := \xi^{GS}_i(\boldsymbol{q}^l)$ for all $\boldsymbol{q} \in \mathcal{L}$, where $\boldsymbol{q}^l \in \{0,1\}$ indicates $Q_l := \{i \in [n] \mid q_i \geq l\}$. Our computed solution is simply the union of the Steiner forests from each round (as implied by ξ^{GS}), where multiple edges count as copies.

Theorem 4. *There is a marginal cost-sharing method χ^{GS} and a solution with cost $C'(\boldsymbol{q})$ for each $\boldsymbol{q} \in \mathcal{Q}$, such that MoulinGD$_{\chi^{GS}}$ is $((2-\frac{1}{n})\cdot\mathcal{H}_L)$-BB, $((2-\frac{1}{n}+\log^2 n)\cdot\mathcal{H}_L)$-EFF, and GSP with respect to C and C'. Furthermore, it satisfies NPT, VP, and CS.*

References

[AKR95] A. Agrawal, P. Klein, and R. Ravi. When trees collide: An approximation algorithm for the generalized steiner problem in networks. *SIAM Journal on Computing*, 24(3):445–456, 1995.

[CRS06] S. Chawla, T. Roughgarden, and M. Sundararajan. Optimal cost-sharing mechanisms for Steiner forest problems. In *Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE'06)*, volume 4286 of *LNCS*, pages 112–123, 2006.

[DMV05] Nikhil R. Devanur, Milena Mihail, and Vijay V. Vazirani. Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decision Support Systems*, 39(1):11–22, 2005.

[GB93] M. Goemans and D. Bertsimas. Survivable networks, linear programming relaxations and the parsimonious property. *Mathematical Programming*, 60:145–166, 1993.

[IMM05] N. Immorlica, M. Mahdian, and V. Mirrokni. Limitations of cross-monotonic cost sharing schemes. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*, pages 602–611, 2005.

[Jai01] K. Jain. A factor 2 approximation algorithm for the generalized steiner network problem. *Combinatorica*, 21(1):39–60, 2001.

[JV01] K. Jain and V. Vazirani. Applications of approximate algorithms to cooperative games. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 364–372, 2001.

[KLS05] J. Könemann, S. Leonardi, and G. Schäfer. A group-strategyproof mechanism for Steiner forests. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*, pages 612–619, 2005.

[Mou99] Hervé Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.

[MRS07] A. Mehta, T. Roughgarden, and M. Sundararajan. Beyond Moulin mechanisms. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC'07)*, pages 1–10, 2007. Full version available at `http://theory.stanford.edu/~tim/papers/bmm.pdf`.

[MS01] Hervé Moulin and Scott Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.

[PT03] M. Pál and É. Tardos. Group strategyproof mechanisms via primal-dual algorithms. In *Proceedings of the 44th Annual Symposium on Foundations of Computer Science (FOCS'03)*, pages 584–593, 2003.

[RS06] T. Roughgarden and M. Sundararajan. New trade-offs in cost-sharing mechanisms. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC'06)*, pages 79–88, 2006.

[RS07] T. Roughgarden and M. Sundararajan. Optimal efficiency guarantees for network design mechanisms. In *Proceedings of the 12th Conference on Integer Programming and Combinatorial Optimization*, pages 469–483, 2007.

[SSL03] D.B. Shmoys, C. Swamy, and R. Levi. Fault-tolerant facility location. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 735–736, 2003.