

Efficient Black-Box Reductions for Separable Cost Sharing

Tobias Harks

Universität Augsburg, Institut für Mathematik, Augsburg, Germany
tobias.harks@math.uni-augsburg.de

Martin Hoefer

Goethe-Universität Frankfurt am Main, Institut für Informatik, Frankfurt am Main, Germany
mhoefer@cs.uni-frankfurt.de

Anja Huber

Universität Augsburg, Institut für Mathematik, Augsburg, Germany
anja.huber@math.uni-augsburg.de

Manuel Surek

Universität Augsburg, Institut für Mathematik, Augsburg, Germany
manuel.surek@math.uni-augsburg.de

Abstract

In *cost sharing games with delays*, a set of agents jointly uses a finite subset of resources. Each resource has a fixed cost that has to be shared by the players, and each agent has a non-shareable player-specific delay for each resource. A prominent example is uncapacitated facility location (UFL), where facilities need to be opened (at a shareable cost) and clients want to connect to opened facilities. Each client pays a cost share and his non-shareable physical connection cost. Given any profile of subsets used by the agents, a *separable cost sharing protocol* determines cost shares that satisfy budget balance on every resource and separability over the resources. Moreover, a separable protocol guarantees existence of pure Nash equilibria in the induced strategic game for the agents.

In this paper, we study separable cost sharing protocols in several general combinatorial domains. We provide black-box reductions to reduce the design of a separable cost sharing protocol to the design of an approximation algorithm for the underlying cost minimization problem. In this way, we obtain new separable cost sharing protocols in games based on arbitrary player-specific matroids, single-source connection games without delays, and connection games on n -series-parallel graphs with delays. All these reductions are efficiently computable – given an initial allocation profile, we obtain a profile of no larger cost and separable cost shares turning the profile into a pure Nash equilibrium. Hence, in these domains any approximation algorithm can be used to obtain a separable cost sharing protocol with a price of stability bounded by the approximation factor.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory, Theory of computation → Exact and approximate computation of equilibria, Theory of computation → Network games

Keywords and phrases Cost Sharing, Price of Stability, Matroids, Connection Games

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.154

Related Version A full version of the paper is available at [34], <https://arxiv.org/abs/1802.10351>.



© Tobias Harks, Martin Hoefer, Anja Huber, and Manuel Surek;
licensed under Creative Commons License CC-BY

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).
Editors: Ioannis Chatzigiannakis, Christos Kaklamani, Dániel Marx, and Donald Sannella;
Article No. 154; pp. 154:1–154:15



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Funding The research of the authors from Augsburg was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - HA 8041/1-1.

1 Introduction

Cost sharing is a fundamental task in networks with strategic agents and has attracted a large amount of interest in algorithmic game theory. Traditionally, cost sharing has been studied in a cooperative sense, i.e., in the form of cooperative games or mechanism design. Many of these approaches treat cost in a *non-separable* way and return a single, global cost share for each agent. In contrast, when agents jointly design a resource infrastructure in large networks, it is much more desirable to provide algorithms and protocols for *separable* cost sharing that specify which agent needs to pay how much to each resource. Here the natural approach are strategic cost sharing games with n players that use subsets of m resources. Each resource generates a cost depending on the subset of players allocating it. A protocol determines a cost share for each resource and each player using it. In addition to separability, there are further natural desiderata for such protocols, such as budget-balance (distribute exactly the arising cost of each resource) and existence of a pure Nash equilibrium (PNE), i.e., allow the resulting game to stabilize.

Perhaps the most prominent such protocol is the fair share protocol, in which the cost of each resource is allocated in equal shares to the players using it. This approach has been studied intensively (see our discussion below), but there are several significant drawbacks. Even in connection games on undirected networks, it can be PLS-hard to find a PNE [51] and the PoS (the total cost of the best Nash equilibrium compared to the cost of the optimal allocation) is not known to be constant. This contrasts the fact that there are polynomial time approximation algorithms with low approximation factors, see, e.g. [12]. For directed networks the PoS can even be as large as $\Omega(\log n)$ [4, 17].

In this paper, we study a slight generalization of cost sharing games, where every resource has a shareable cost component and a non-shareable player-specific delay component. The shareable cost needs to be shared by the players using it, the non-shareable player-specific delay represents, e.g., a physical delay and is thus unavoidable. This setting arises in several relevant scenarios, such as uncapacitated facility location (UFL) [37]. Here players share the monetary cost of opened facilities but additionally experience delays measured by the distance to the closest open facility. Another important example appears in network design, where players jointly buy edges of a graph to connect their terminals. Besides the monetary cost for buying edges, each player experiences player-specific delays on the chosen paths. In such a distributed network environment, it is not clear a priori if an optimal solution can be stable – i.e., if the shareable costs can be distributed among the players in a separable way so that players do not want to deviate from it. This question leads directly to the design of protocols that distribute the costs in order to induce stable and good-quality solutions of the resulting strategic game.

Our results are three polynomial-time black-box reductions for the price of stability (PoS) of separable cost sharing protocols in combinatorial resource allocation problems. Our domains represent broad generalizations of UFL – arbitrary, player-specific matroids, single-source connection games without delays, and connection games on undirected n -series-parallel graphs with delays. In each of these domains, we take as input an arbitrary profile and efficiently turn it into a profile having no larger cost and a sharing of the shareable costs such that it is a Nash equilibrium. Our protocols are polynomial-time in several ways. Firstly, the games we study are succinctly represented. In matroids, we assume that strategies

are represented implicitly via an independence oracle. For connection games on graphs, the strategies of each player are a set of paths, which is implicitly specified by terminal vertices of the player and the graph structure. The cost sharing protocol is represented by a strategy profile S and a sharing of the shareable costs arising in S on each resource. While in principle the protocol must specify a sharing of the costs for all of the other (possibly exponentially many) strategy profiles, one can do so implicitly by a simple lexicographic assignment rule. It guarantees that the profile S becomes a PNE. As such, starting from an arbitrary initial profile S' , we can give in polynomial time the Nash equilibrium profile S , the cost shares for S , and the assignment rule for cost shares in the other profiles. Hence, if S' is polynomial-time computable, then both protocol and Nash equilibrium S are both polynomial-time computable and polynomial-space representable.

1.1 Our Results

We present several new polynomial-time black-box reductions for separable cost sharing protocols with small PoS. We study three domains that represent broad generalizations of the uncapacitated facility location problem. In each domain, we devise an efficient black-box reduction that takes as input an arbitrary strategy profile and computes a new profile of no larger cost together with a separable cost sharing protocol inducing the new profile as a PNE. Thus, *any* polynomial-time α -approximation of the social cost can be turned into a separable cost sharing protocol with PoS at most α .

Matroidal Set Systems. In Section 3 we provide a black-box reduction for matroidal set systems. Our results even apply to the broad class of *subadditive* cost functions that include fixed costs and discrete concave costs even with weighted players as a special case. Here we assume access to a value oracle for the subadditive cost function for each resource. Matroidal set systems with player-specific delays include uncapacitated facility location as a special case, since these correspond to matroid games, where each player has a uniform rank 1 matroid. For *metric* UFL, there is for instance a 1.488-approximation algorithm [45] using ideas of a previous 1.5-approximation algorithm [11]. This leads to a separable cost sharing protocol with PoS equal 1.488. Also, the existing hardness results for UFL carry over, meaning that for any polynomial-time computable, separable cost sharing protocol, the cost of an equilibrium that can be computed in polynomial time is bounded from below by the inapproximability bound. For metric UFL there is a lower bound of 1.46 [31].

Single-Source Connection Games with Fixed Costs. In Section 4 we consider cost sharing games on graphs, where the set systems correspond to paths connecting a global source with a player-specific terminal. We again provide a polynomial-time black-box reduction. Our result improves significantly over the existing Prim-Sharing [17] with a PoS of 2. We obtain separable protocols based on any approximation algorithm for Steiner tree, such as, e.g., the classic 1.55-approximation algorithm [49], or the celebrated recent 1.39-approximation algorithm [12]. Our black-box reduction continues to hold even for directed graphs, where we can use any algorithm for the Directed Steiner Tree problem [15], or games based on the (directed or undirected) Group Steiner Tree problem [16, 26]. All lower bounds on approximation hardness translate to the quality of polynomial-time computable equilibria of polynomial-time computable separable protocols.

Connection Games With Delays. Finally, in Section 5 we study multi-terminal connection games with delays and fixed costs. For directed graphs, an optimal Steiner forest is not enforceable by a separable cost sharing protocol, even for two players [17]. Very recently,

a similar result was shown even for two-player games on undirected graphs [35]. Thus, for general graphs, we cannot expect separable protocols with optimal or close-to-optimal equilibria, or (efficient) black-box reductions. We introduce a class of so-called n -series-parallel graphs, which allows to obtain a black-box reduction in polynomial time. The transformation directly implies that the n -series-parallel graphs always admit a separable cost sharing protocol inducing an optimal Steiner forest as an equilibrium.

The reduction also applies to discrete-concave cost functions and player-specific delays, however, we do not know if polynomial running time is guaranteed. n -series-parallel graphs have treewidth at most 2, thus, for fixed edge costs and no delays, it is possible to compute efficiently even an optimal Steiner forest [7]. Hence, in this case we obtain a separable protocol with PoS of 1 in polynomial time. We finally demonstrate that the specific setting of n -series-parallel graphs is in some sense necessary: Even for generalized series-parallel graphs we give a counterexample showing that a black-box reduction is impossible to achieve.

1.2 Preliminaries and Related Work

Cooperative cost sharing games have been studied over the last decades for a variety of combinatorial optimization problems, such as minimum spanning tree [10], Steiner tree [29,30,46,52], facility location [28], vertex cover [22], and many more. Cooperative cost sharing games have interesting implications for (group-)strategyproof cost sharing mechanisms [41,42,47,48]. For Bayesian cost-sharing mechanisms there even exist efficient black-box reductions from algorithm to mechanism design [27]. A major difference to our work is that cooperative cost sharing is *not separable*.

The most prominent example of a *separable* cost sharing protocol is the *fair share protocol*, in which the cost of each resource is divided in equal shares among the players that use it. This protocol is also anonymous, and it implies that the resulting game is a congestion game [50]. It guarantees the smallest PoS within a class of anonymous protocols [17]. The fair share protocol has attracted a serious amount of research interest over the last decade [1,4,8,33], especially the notorious open problem of a constant PoS for connection games in undirected graphs [9,23,25,43,44]. However, as a significant drawback - outside of the domain of undirected connection games - the PoS is often as large as $\Omega(\log n)$, e.g. [17,32].

More general separable protocols have been studied mostly in terms of the price of anarchy, e.g., for scheduling (or matroid games) [6,13,19,24,53] or single-source network design with [20,21] and without uncertainty [17]. The best result here is a price of anarchy (and stability) of 2 via Prim-Sharing [17], a protocol inspired by Prim's MST algorithm. A protocol with logarithmic PoS was shown for capacitated UFL games [37].

We observe that separable protocols with low PoS can be obtained using results for cost sharing games with so-called "arbitrary sharing", where effectively players pay the cost increment when changing their strategy (see [36] for a formal definition). A PNE for arbitrary sharing can be translated directly into a PNE for a suitable separable cost sharing protocol. A simple proof of Proposition 1 can be found in the full version of this paper [34].

► **Proposition 1.** *If for a cost sharing model, the non-cooperative game with arbitrary sharing has a PNE, then there is a separable cost sharing protocol with the same PNE.*

This implies existence of separable protocols with optimal PNE and PoS 1 for a variety of classes of games, including matroid games with uniform discrete-concave costs [36], uncapacitated facility location with fixed [14] and discrete-concave costs [39], connection games (single-source [5,38] and other classes [2,3,40]) with fixed costs, and more. However, the large majority of these results are *inefficient*, i.e., there is no polynomial-time algorithm that computes the required optimal equilibrium.

Alternatively, one may resort to approximate equilibria in games with arbitrary sharing that are efficiently computable. The most prominent technique works via reducing costs by an additive value ε to ensure polynomial running time (put forward for single-source connection games in [5] and used in much of the follow-up work [2, 3, 14, 38]). This approach *does not translate* to separable protocols, since a player must eventually contribute to *all resources*. This is impossible for the model we consider here.

2 Separable Cost Sharing Protocols

We are given a finite set N of players and a finite set E of resources. Each player $i \in N$ is associated with a predefined family of subsets $\mathcal{S}_i \subseteq 2^E$ from which player i needs to pick at least one. The space of strategy profiles is denoted by $\mathcal{S} := \times_{i \in N} \mathcal{S}_i$. For $S \in \mathcal{S}$ we denote by $N_e(S) = \{i \in N : e \in S_i\}$ the set of players that use resource e . Every resource $e \in E$ has a fixed cost $c_e \geq 0, e \in E$ that is assumed to be *shareable* by the players. In addition to the shareable costs, there are *player-specific constant costs* $d_{i,e} \geq 0, i \in N, e \in E$ that are not shareable. If player i chooses subset S_i , then the player-specific costs $\sum_{e \in S_i} d_{i,e}$ must be paid completely by player i . The total cost of a profile S is defined as $C(S) = \sum_{e \in \cup_{i \in N} S_i} c_e + \sum_{i \in N} \sum_{e \in S_i} d_{i,e}$.

A *cost sharing protocol* assigns cost share functions $\xi_{i,e} : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ for all $i \in N$ and $e \in E$ and thus induces the strategic game (N, \mathcal{S}, ξ) . For a player i , her total private cost of strategy S_i in profile S is $\pi_i(S) := \sum_{e \in S_i} (\xi_{i,e}(S) + d_{i,e})$. We assume that every player picks a strategy in order to minimize her private cost. A prominent solution concept in non-cooperative game theory are pure Nash equilibria. Using standard notation in game theory, for a strategy profile $S \in \mathcal{S}$ we denote by $(S'_i, S_{-i}) := (S_1, \dots, S_{i-1}, S'_i, S_{i+1}, \dots, S_n) \in \mathcal{S}$ the profile that arises if only player i deviates to strategy $S'_i \in \mathcal{S}_i$. A profile is a *pure Nash equilibrium (PNE)* if for all $i \in N$ it holds $\pi_i(S) \leq \pi_i(S'_i, S_{-i})$ for all $S'_i \in \mathcal{S}_i$.

In order to be practically relevant, cost sharing protocols need to satisfy several desiderata. In this regard, *separable* cost sharing protocols are defined as follows [17].

► **Definition 2** (Cost Sharing Protocols and Enforceability). A cost sharing protocol is

1. *stable* if it induces only games that admit at least one pure Nash equilibrium.
2. *budget balanced*, if for all $e \in E$ with $N_e(S) \neq \emptyset$

$$c_e = \sum_{i \in N_e(S)} \xi_{i,e}(S) \text{ and } \xi_{i,e}(S) = 0 \text{ for all } i \notin N_e(S).$$

3. *separable* if it is stable, budget-balanced and induces only games for which in any two profiles $S, S' \in \mathcal{S}$ for every resource $e \in E$,

$$N_e(S) = N_e(S') \Rightarrow \xi_{i,e}(S) = \xi_{i,e}(S') \text{ for all } i \in N_e(S).$$

4. *polynomial time computable*, if the cost sharing functions ξ can be computed in polynomial time in the encoding length of the cost sharing game.

We call a strategy profile S *enforceable*, if there is a separable protocol inducing S as a PNE.

Separability means that for any two profiles S, S' the cost shares on e are the same if the set of players using e remains unchanged. Still, separable protocols can assign cost share functions that are specifically tailored to a given congestion model, for example based on an optimal profile. In this paper, we are additionally interested in *polynomial-time computable protocols* that we introduce here.

3 Matroid Games

In this section, we consider matroid games. As usual in matroid theory, we will write \mathcal{B}_i instead of \mathcal{S}_i , and \mathcal{B} instead of \mathcal{S} , when considering matroid games. The tuple $\mathcal{M} = (N, E, \mathcal{B}, (c_e)_{e \in E}, (d_{i,e})_{e \in E, i \in N})$ is called a *matroid game* if $E = \bigcup_{i \in N} E_i$, and each set system $\mathcal{B}_i \subseteq 2^{E_i}$ forms the base set of some matroid $\mathcal{M}_i = (E_i, \mathcal{B}_i)$. While seemingly abstract, the class includes several prominent application domains, such as UFL games. In a UFL game, the resources are facilities (e.g. common transport hubs) and the players incur delay $d_{i,e}$ in addition to their cost shares for opening used facilities. Every player i chooses exactly one resource, that is $|B_i| = 1$ for all $B_i \in \mathcal{B}_i$ and $i \in N$ and hence \mathcal{B}_i corresponds to a *uniform matroid of rank one*. Recall that every base B of a matroid $\mathcal{M}_i = (E_i, \mathcal{B}_i)$ has the same cardinality which we denote with rk_i (the rank of \mathcal{M}_i).

In the following, instead of fixed costs on the resources, we allow for *subadditive* cost functions $c_e : 2^N \rightarrow \mathbb{R}_{\geq 0}, e \in E$. c_e is called *subadditive*, if it satisfies (1) $c_e(S) \leq c_e(T)$ for all $S \subseteq T \subseteq N$, and (2) $c_e(S + \{i\}) \leq c_e(S) + c_e(\{i\})$ for all $S \subset N, i \in N$. Note that subadditive functions include fixed costs, but also discrete concave costs, where the nondecreasing cost $c_e : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ depends on the number of players using e and satisfies $c_e(x + \delta) - c_e(x) \geq c_e(y + \delta) - c_e(y) \forall x \leq y; x, y, \delta \in \mathbb{N}$. As such, in the discrete concave setting, all players have the same weight of 1, whereas subadditive costs also allow for different weights (as in weighted congestion games). We furthermore assume $c_e(\emptyset) = 0, e \in E$.

Let us denote the cost of the cheapest alternative of player i to resource e for profile $B \in \mathcal{B}$ by $\Delta_i^e(B) := \min_{f \in E, B_i + f - e \in \mathcal{B}_i} (c_e(B_i + f - e, B_{-i}) + d_{i,f})$. Here we use the simplified notations $B_i + f - e := B_i \cup \{f\} \setminus \{e\}$ and $c_e(B) := c_e(N_e(B))$. We recapitulate the following characterization of enforceable strategy profiles (obtained in [37], where also a cost sharing protocol inducing the enforceable profile as a PNE is given).¹

► **Lemma 3.** *A profile $B = (B_1, \dots, B_n)$ is enforceable iff the following two properties are satisfied. Note that (D1) implies that each summand $\Delta_i^e(B) - d_{i,e}$ in (D2) is nonnegative.*

$$d_{i,e} \leq \Delta_i^e(B) \text{ for all } i \in N, e \in B_i \quad (\text{D1})$$

$$c_e(B) \leq \sum_{i \in N_e(B)} (\Delta_i^e(B) - d_{i,e}) \text{ for all } e \in E. \quad (\text{D2})$$

► **Remark.** The characterization was used in [37] to prove that an optimal collection of bases is enforceable. This implies a PoS of 1 for a separable cost sharing protocol that relies on the optimal profile. As such, the protocol is not efficiently computable (unless $P = NP$).

In the following, we devise a black-box reduction in Algorithm 1. It takes as input an arbitrary collection of bases B and transforms them *in polynomial time* into an enforceable set of bases B' of no larger cost. We define for each $i \in N, e \in E$ a *virtual cost value* $v_i^e := c_e(\{i\}) + d_{i,e}$, and for each $B \in \mathcal{B}, i \in N, e \in E$ a *virtual deviation cost* $\bar{\Delta}_i^e(B) := \min_{f \in E, B_i + f - e \in \mathcal{B}_i} v_i^f$. The algorithm now iteratively checks whether (D1) and (D2) from Lemma 3 hold true (in fact it checks this condition for smaller values on the right hand side given by the virtual values), and if not, exchanges one element of some player. We show that the algorithm terminates with an enforceable profile after polynomially many steps.

► **Theorem 4.** *Let B be a strategy profile for a matroid congestion model with subadditive costs. There is an enforceable profile B' with $C(B') \leq C(B)$ that can be computed in at most*

¹The original characterization in [37] was proven for weighted players and load-dependent non-decreasing cost functions but the proof also works for subadditive cost functions.

Algorithm 1: Transforming any profile B into an enforceable profile B'

Input: Congestion model (N, \mathcal{B}, c, d) and profile $B \in \mathcal{B}$ **Output:** Enforceable profile B' with $C(B') \leq C(B)$.

```

1 Set  $B' \leftarrow B$ 
2 while there is  $e \in E$  that satisfies at least one of the following conditions:
     $d_{i,e} > \bar{\Delta}_i^e(B')$  for some  $i \in N_e(B')$  (1) or
     $c_e(B') > \sum_{i \in N_e(B')} (\bar{\Delta}_i^e(B') - d_{i,e})$  (2) do
3     if (1) holds true for some  $i \in N_e(B')$  then
4         Let  $f_i \in \arg \min_{f \in E} v_i^f$ 
5         Update  $B'_i \leftarrow B'_i + f_i - e$ 
6     else if (2) holds true then
7         while (2) holds true on  $e$  do
8             Pick  $i \in N_e(B')$  with  $v_i^e > \bar{\Delta}_i^e(B')$ 
9             Let  $f_i \in \arg \min_{f \in E} v_i^f$ 
10            Update  $B'_i \leftarrow B'_i + f_i - e$ 

```

$n \cdot m \cdot rk(\mathcal{B})$ iterations of the outer while-loop in Algorithm 1, where $rk(\mathcal{B}) = \max_{i \in N} rk_i$. Furthermore, each such iteration needs $\mathcal{O}(n \cdot m \cdot rk(\mathcal{B}) \cdot Q + (m + n) \cdot Q')$ time, where Q and Q' denote the maximum complexities of the independence oracles for the players' strategies, and of the value oracles for the subadditive cost functions of the resources, respectively.

Proof. First, observe that if (D1) and (D2) from Lemma 3 hold true for smaller values $0 \leq \bar{\Delta}_i^e(B) \leq \Delta_i^e(B), i \in N, e \in E$, then the profile B is also enforceable. Hence, if the algorithm terminates, the resulting strategy profile B' will be enforceable.

To show that the algorithm is well-defined, we only need to check Line 8. By subadditivity we get $\sum_{i \in N_e(B')} c_e(\{i\}) > c_e(B')$. Thus, whenever $c_e(B') > \sum_{i \in N_e(B')} (\bar{\Delta}_i^e(B') - d_{i,e})$, there is an $i \in N_e(B')$ with $c_e(\{i\}) + d_{i,e} > \bar{\Delta}_i^e(B')$.

We now bound the running time. Consider player i and the matroid bases \mathcal{B}_i . We interpret a basis $B_i \in \mathcal{B}_i$ as distributing exactly rk_i unit sized packets over the resources in E . This way, we can interpret the algorithm as iteratively moving packets away from those resources $e \in E$ for which either (1) or (2) holds true. We give each packet a unique ID $i_k, k = 1, \dots, rk_i$. For $B_i \in \mathcal{B}_i$, let e_{i_k} denote the resource on which packet i_k is located. We now analyse the two types of packet movements during the execution of the algorithm. For a packet movement executed in Line 5 of Algorithm 1, we have $d_{i,e} > \bar{\Delta}_i^e(B')$, thus, when packet i_k located on $e = e_{i_k}$ is moved to f_i , it holds that $v_i^{e_{i_k}} = v_i^e \geq d_{i,e} > \bar{\Delta}_i^e(B') = v_i^{f_i}$. For packet movements executed in Line 10, then by the choice of player $i \in N_e(B')$ (see Line 8) for the corresponding packet i_k it holds $v_i^{e_{i_k}} = v_i^e > \bar{\Delta}_i^e(B') = v_i^{f_i}$. In both cases we obtain $v_i^e > v_i^{f_i}$. Hence, every movement of a single packet i_k is in strictly decreasing order of virtual value of the resource. Note that the virtual cost value v_i^e does not depend on profile B . Thus, there are at most m different virtual cost values that a packet i_k of player i can experience, and thus packet i_k can move at most $m - 1$ times. The following is an upper bound on the

total number of packet movements for all players $\sum_{i \in N} \text{rk}_i \cdot (m-1) \leq n \cdot m \cdot \text{rk}(\mathcal{B})$. Thus, the number of iterations of the outer while-loop can be at most that value. It is straightforward to observe the stated complexity of one iteration of the while-loop.

It is left to argue that the final output B' has cost at most $C(B)$. We prove this inductively by the different types of packet movements. Consider first a packet movement of type (1). Let B and B' be the profiles before and after packet i_k has been moved from e to f_i , respectively. We obtain (by using subadditivity and assumption (1))

$$\begin{aligned} C(B') - C(B) &= (c_{f_i}(B') - c_{f_i}(B) + d_{i,f_i}) - (c_e(B) - c_e(B') + d_{i,e}) \\ &\leq c_{f_i}(\{i\}) + d_{i,f_i} - (c_e(B) - c_e(B') + d_{i,e}) \\ &= \bar{\Delta}_i^e(B) - d_{i,e} + (c_e(B') - c_e(B)) \leq \bar{\Delta}_i^e(B) - d_{i,e} < 0. \end{aligned}$$

Now consider packet movements of type (2). We treat all movements occurring in one run of the while loop in Line 7. Let B denote the profile before and B' after all these movements. Let $T_e(B) \subseteq N_e(B)$ denote the set of those players whose packet i_k on e is moved to f_i for $i \in T_e(B)$ during the while loop. Let $F_e(B) = \bigcup_{i \in T_e(B)} \{f_i\}$ and for $i \in T_e(B)$ define $T_{f_i}(B) = \{j \in T_e(B) \mid f_j = f_i\}$. We derive some useful observations. Before entering the while loop, it holds

$$c_e(B) > \sum_{i \in N_e(B)} (\bar{\Delta}_i^e(B) - d_{i,e}) = \sum_{i \in N_e(B) \setminus T_e(B)} (\bar{\Delta}_i^e(B) - d_{i,e}) + \sum_{i \in T_e(B)} (\bar{\Delta}_i^e(B) - d_{i,e}). \quad (3)$$

Moreover, after exiting the while loop it holds

$$c_e(B') \leq \sum_{i \in N_e(B) \setminus T_e(B)} (\bar{\Delta}_i^e(B) - d_{i,e}). \quad (4)$$

Thus, combining (3) and (4) we get

$$c_e(B) - c_e(B') > \sum_{i \in T_e(B)} (\bar{\Delta}_i^e(B) - d_{i,e}). \quad (5)$$

Putting everything together, we obtain

$$\begin{aligned} C(B') - C(B) &= \sum_{f_i \in F_e(B)} (c_{f_i}(B') - c_{f_i}(B)) + \sum_{i \in T_e(B)} d_{i,f_i} - (c_e(B) - c_e(B') + \sum_{i \in T_e(B)} d_{i,e}) \\ &\leq \sum_{f_i \in F_e(B)} \sum_{j \in T_{f_i}(B)} c_{f_i}(\{i\}) + \sum_{i \in T_e(B)} d_{i,f_i} - (c_e(B) - c_e(B') + \sum_{i \in T_e(B)} d_{i,e}) \\ &= \sum_{i \in T_e(B)} \bar{\Delta}_i^e(B) - (c_e(B) - c_e(B') + \sum_{i \in T_e(B)} d_{i,e}) < 0, \end{aligned}$$

where the first inequality follows from subadditivity, and the last inequality from (5). ◀

4 Single-Source Connection Games without Delays

In this section, we study connection games in an undirected graph $G = (V, E)$ with a common source vertex $s \in V$. Every player i wants to connect a player-specific terminal node $t_i \in V$ to s . Consequently, every strategy P_i of player i is an (s, t_i) -path in G . We denote the set of paths for player i by \mathcal{P}_i and the set of profiles by \mathcal{P} . Furthermore we focus on fixed shareable costs $c_e \geq 0$ and no player-specific delays $d_{i,e} = 0$, for all $i \in N$, $e \in E$.

For single-source games with delays, a black-box reduction is impossible to achieve, since one can construct instances where no optimal solution is enforceable. To see this, take any

multi-source multi-terminal connection game and introduce a new auxiliary source vertex s . Then connect s to each s_i with an auxiliary edge e_i , which has cost $c_{e_i} = 0$ and delays $d_{i,e_i} = 0$, $d_{j,e_i} = M$ for all other players j (for some prohibitively large constant M). Now in any equilibrium and any optimal state of the resulting game, player i will choose an (s, t_i) -path which begins with edge e_i . Moreover, e_i does not generate additional cost for player i . As such, the optimal solutions, the Nash equilibria, and their total costs correspond exactly to the ones of the original multi-source game. For the general multi-source case, there are instances with no enforceable optimal solution (even for only two players, see [35]).

For connection games with fixed shareable costs and no delays, the state of the art for polynomial-time cost sharing protocols is the Prim-Sharing protocol with a PoS of 2 [17]. Moreover, it is known that an optimal *tree* profile (a profile in which the union of player paths constitute a tree) is always enforceable [5,17]. Our result for polynomial-time computation of cheap enforceable profiles, stated in Theorem 5, represents a significant generalization – for any (possibly non-enforceable) tree profile, Algorithm 2 (given in detail in [34]) computes an enforceable profile that can only be cheaper in terms of total cost. Moreover, by combining the result with existing approximation algorithms, we obtain a protocol with PoS of 1.39.

► **Theorem 5.** *Let P be a strategy profile for a single-source connection game with fixed costs. There is an enforceable profile P' with $C(P') \leq C(P)$ that can be computed by Algorithm 2 in polynomial time.*

It is straightforward that for fixed costs we can transform each profile P into a tree profile \hat{P} , in which the union of player paths constitute a tree T , without increasing the cost. Over the course of the algorithm, we adjust this tree and construct a cost sharing for it in a bottom-up fashion. The approach has similarities to an approach for obtaining approximate equilibria for single-source cost sharing games with arbitrary sharing [5]. However, our algorithm exploits crucial properties of separable protocols, thereby providing an exact Nash equilibrium and polynomial running time.

When designing a separable protocol based on a state \hat{P} , we can always assume that when a player i deviates unilaterally to one or more edges $e \in G \setminus \hat{P}_i$, she needs to pay all of c_e . As such, player i always picks a collection of shortest paths with respect to c_e between pairs of nodes on her current path \hat{P}_i . All these paths in G are concisely represented in the algorithm as “auxiliary edges”. The algorithm initially sets up an auxiliary graph \hat{G} given by T and the set of auxiliary edges based on \hat{P} . It adjusts the tree T by removing edges of T and adding auxiliary edges in a structured fashion.

We first show in the following Lemma 6 (for a proof see [34]) that this adjustment procedure improves the total cost of the tree, and that the final tree \hat{T} is enforceable in \hat{G} . In the corresponding cost sharing, every auxiliary edge contained in \hat{T} is completely paid for by a single player that uses it. In the subsequent proof of Theorem 5, we only need to show that for the auxiliary edges in \hat{T} , the edge costs of the corresponding shortest paths in G can be assigned to the players such that we obtain an equilibrium in G . The proof shows that the profile P' evolving in this way is enforceable in G and has cost no larger than P .

► **Lemma 6.** *Algorithm 2 computes a cost sharing of a feasible tree \hat{T} in the graph \hat{G} . The total cost $C(\hat{T}) \leq C(T)$, every auxiliary edge in \hat{T} is paid for by a single player, and the corresponding profile \hat{P} is enforceable in \hat{G} .*

Proof of Theorem 5. The previous lemma shows that the algorithm computes a cost sharing of a tree \hat{T} in \hat{G} , such that every player is happy with the path \hat{P}_i and every auxiliary edge in \hat{T} is paid for completely by a single player. We now transform \hat{P} into P' by replacing each

auxiliary edge $e = (u, v) \in \hat{P}_i$ by the corresponding shortest path $P(u, v)$ in G . We denote by E_i the set of edges introduced in the shortest paths for auxiliary edges in \hat{P}_i . For the total cost of the resulting profile we have that $C(P') \leq C(\hat{P}) \leq C(P)$, since the sets E_i can overlap with each other or the non-auxiliary edges of \hat{T} .

We show that P' is enforceable by transforming the cost sharing constructed in function \hat{c} into separable cost sharing functions as follows. Initially, set $\xi_{i,e}(P') = 0$ for all $e \in E$ and $i \in N$. Then, for each non-auxiliary edge $e \in \hat{T}$ we assign $\xi_{i,e}(P') = \hat{c}_e(i)$ if $e \in \hat{P}_i$ and $\xi_{i,e}(P') = 0$ otherwise. Finally, number players arbitrarily from 1 to n and proceed in that order. For player i , consider the edges in E_i . For every $e \in E_i$, if $\sum_{j < i} \xi_{j,e}(P') = 0$, then set $\xi_{i,e}(P') = c_e$.

This yields a budget-balanced assignment for state P' . As usual, if a player i deviates in P' from P'_i to P''_i , we can assume player i is assigned to pay the full cost c_e for every edge $e \in P''_i \setminus P'_i$. To show that there is no profitable deviation from P' , we first consider a thought experiment, where every edge in E_i comes as a separate edge bought by player i . Then, clearly P' is enforceable – the cost of P'_i with ξ is exactly the same as the cost of \hat{P} with \hat{c} in \hat{G} . Moreover, any deviation P''_i can be interpreted as an (s, t_i) -path in \hat{G} by replacing all subpaths consisting of non-auxiliary edges in P'' by the corresponding auxiliary edge of \hat{G} . As such, the cost of P''_i is exactly the same as the cost of the corresponding deviation in \hat{G} . Now, there is not a separate copy for every edge in E_i . The set E_i can overlap with other sets E_j and/or non-auxiliary edges. Then player i might not need to pay the full cost on some $e \in E_i$. Note, however, every edge for which player i pays less than c_e is present in P'_i as well. Hence, P''_i cannot improve over P'_i due to this property. ◀

The result continues to hold for various generalizations. For example, we can immediately apply the arguments in directed graphs, where every player i seeks to establish a directed path between t_i and s . Moreover, the proof can also be applied readily for a group-connection game, where each player wants to establish a directed path to s from *at least one node of a set* $V_i \subset V$. For this game, we simply add a separate super-terminal t_i for every player i and draw a directed edge of cost 0 from t_i to every node in V_i .

► **Corollary 7.** *Let P be a strategy profile for a single-source group-connection game in directed graphs with fixed costs. There is an enforceable profile P' with $C(P') \leq C(P)$ that can be computed by Algorithm 2 in polynomial time.*

5 Connection Games and Graph Structure

In this section, we consider connection games played in undirected graphs $G = (V, E)$ with player-specific source-terminal pairs. Each player $i \in N$ has a source-terminal-pair (s_i, t_i) . Consequently, every strategy P_i of player i is an (s_i, t_i) -path in G . We denote the set of paths for player i by \mathcal{P}_i .

Note that we can assume w.l.o.g. that $(G, (s_1, t_1), \dots, (s_n, t_n))$ is *irredundant*, meaning that each edge and each vertex of G is contained in at least one (s_i, t_i) -path for some player $i \in N$ (nodes and edges not used by any player can easily be recognized (and then deleted) by Algorithm IRREDUNDANT in [34]; adapted from Algorithm 1 in [18]).

For the special case without delays, enforceability was characterized via an LP in [35]. We

can directly adapt this characterization as follows. For a strategy profile $P = (P_1, \dots, P_n)$:

$$\begin{aligned} \text{LP}(P) \quad \max \quad & \sum_{i \in N, e \in P_i} \xi_{i,e} \quad \text{s.t.:} \quad \sum_{i \in N_e(P)} \xi_{i,e} \leq c_e \quad \forall e \in E \text{ with } N_e(P) \neq \emptyset, \\ & \sum_{e \in P_i \setminus P'_i} (\xi_{i,e} + d_{i,e}) \leq \sum_{e \in P'_i \setminus P_i} (c_e + d_{i,e}) \quad \forall P'_i \in \mathcal{P}_i \quad \forall i \in N \quad (\text{NE}) \\ & \xi_{i,e} \geq 0 \quad \forall e \in P_i \quad \forall i \in N \end{aligned}$$

► **Theorem 8.** P is enforceable iff there is an optimal solution $(\xi_{i,e})_{i \in N, e \in P_i}$ for $\text{LP}(P)$ with $\sum_{i \in N_e(P)} \xi_{i,e} = c_e \quad \forall e \in E \text{ with } N_e(P) \neq \emptyset.$ (BB)

Given an optimal solution $(\xi_{i,e})_{i \in N, e \in P_i}$ for $\text{LP}(P)$ with the property (BB), the profile P becomes a PNE in the game induced by ξ , which assigns for each $i \in N$ and $e \in E$ and each strategy profile $P' = (P'_1, \dots, P'_n)$ the following cost shares (these cost shares resemble those introduced in [53]; the proof that P is a PNE can be directly adapted from [35]):

$$\xi_{i,e}(P') = \begin{cases} \xi_{i,e}, & \text{if } i \in S_e(P') = S_e(P), \\ c_e, & \text{if } i \in (S_e(P') \setminus S_e(P)) \text{ and } i = \min(S_e(P') \setminus S_e(P)), \\ c_e, & \text{if } i \in S_e(P') \subsetneq S_e(P) \text{ and } i = \min S_e(P'), \\ 0, & \text{else.} \end{cases}$$


We now introduce a subclass of generalized series-parallel graphs for which we design a polynomial time black-box reduction.

► **Definition 9** (n -series-parallel graph). An irredundant graph $(G, (s_1, t_1), \dots, (s_n, t_n))$ is n -series-parallel if, for all $i \in N$, the subgraph G_i (induced by \mathcal{P}_i) is created by a sequence of series and/or parallel operations starting from the edge $s_i - t_i$. For an edge $e = u - v$, a series operation replaces it by a new vertex w and two edges $u - w, w - v$; a parallel operation adds to $e = u - v$ a parallel edge $e' = u - v$.

► **Theorem 10.** If $(G, (s_1, t_1), \dots, (s_n, t_n))$ is n -series-parallel, the following holds:

- (1) Given an arbitrary strategy profile P , an enforceable strategy profile P' with cost $C(P') \leq C(P)$, and corresponding cost share functions ξ , can be computed in polynomial time.
- (2) For all cost functions c, d , every optimal strategy profile of $(G, (s_1, t_1), \dots, (s_n, t_n), c, d)$ is enforceable.
- (3) For all edge costs c , an optimal Steiner forest of $(G, (s_1, t_1), \dots, (s_n, t_n), c)$ can be computed in polynomial time.

Proof Sketch for Theorem 10. We first describe how to compute, given an arbitrary profile $P = (P_1, \dots, P_n)$, an enforceable strategy profile with cost at most $C(P)$. Assume that P is not enforceable, and let $(\xi_{i,e})_{i \in N, e \in P_i}$ be an optimal solution for $\text{LP}(P)$. In the following, we denote the variables $(\xi_{i,e})_{i \in N, e \in P_i}$ as *cost shares*, although they do not correspond to a budget-balanced cost sharing protocol (since P is not enforceable). There is at least one edge f which is not completely paid, i.e. $\sum_{i \in N_f(P)} \xi_{i,f} < c_f$ holds. The optimality of the cost shares implies that each player $i \in N_f(P)$ has an alternative path P'_i with $f \notin P'_i$ and equality in the corresponding $\text{LP}(P)$ -inequality (so-called *tight alternative* of player i that substitutes f). Furthermore, if the path P_i of player i contains more than one edge which is not completely paid, there is a tight alternative P'_i for player i which substitutes all non-paid edges. Figure 1 illustrates this for the case that P_i (straight edges) contains the three non-paid edges f, g, h , and player i substitutes them by using P'_i (thick edges).



■ **Figure 1** Illustration for substitution of nonpaid edges.

The high-level idea of the algorithm now is that all players with unpaid edges in their paths deviate from those edges by using tight alternatives until we reach a strategy profile in which all edges are completely paid. This strategy profile will then be enforceable and cheaper than P . In the following we explain this in more detail. Let $P' = (P'_1, \dots, P'_n)$ be the strategy profile which results from P if all players with unpaid edges in their paths substitute all these edges by a tight alternative path (as described above). Furthermore we define cost shares (again not nec. budget-balanced) for P' (for each player $i \in N, e \in P'_i$):

$$\xi_{i,e}(P') = \begin{cases} \xi_{i,e}, & \text{for } e \in P'_i \cap P_i, \\ c_e, & \text{for } e \in P'_i \setminus P_i. \end{cases}$$

If $\sum_{i \in N_e(P')} \xi_{i,e}(P') \geq c_e$ holds for all edges e with $N_e(P') \neq \emptyset$, the profile P' has the desired properties: The cost of P' is strictly smaller than P since the private costs of the players remain unchanged if they use tight alternatives. Furthermore the cost shares $(\xi_{i,e}(P'))_{i \in N, e \in P'_i}$ induce a feasible solution of $\text{LP}(P')$ with (BB); thus P' is enforceable (we possibly need to decrease some cost shares to get a feasible solution of $\text{LP}(P')$).

It remains to consider the case that there is at least one edge f which is not completely paid, i.e. for which $\sum_{i \in N_f(P')} \xi_{i,f}(P') < c_f$ holds. We can show that all users of nonpaid edges again have tight alternatives which substitute those edges, therefore we can again update the strategy profile (resulting in P'') and the corresponding cost shares. If now all edges are completely paid, P'' has the desired properties. Proceeding in this manner, we can show that one finally reaches a strategy profile for which all edges are completely paid; thus it is enforceable and cheaper than the profile P . Algorithm n -SEPA in [34] summarizes the described procedure. To complete the proof of the first statement of Theorem 10, it remains to show that P' and ξ can be computed in polynomial time, i.e. Algorithm n -SEPA has polynomial running time. This follows from two facts. First, the number of strategy profiles that we have to consider until we reach the desired strategy profile is bounded by $|P|$, the number of edges in the union of the paths P_1, \dots, P_n . Second, we can solve $\text{LP}(P)$ in polynomial time. To this end we show that, for every player i , we do not need to consider all paths $P'_i \in \mathcal{P}_i$ in (NE) of $\text{LP}(P)$, which can be exponentially many paths, but only a set of alternatives \mathcal{A}_i of polynomial cardinality. Algorithm $\text{ALTERNATIVES}(i)$ in [34] computes this set of alternatives. This completes the proof sketch for the first statement of Theorem 10.

Our analysis of the cost of P' immediately implies that every optimal strategy profile has to be enforceable: Otherwise Algorithm n -SEPA computes a strategy profile with strictly smaller cost; contradiction. Thus the second statement of Theorem 10 holds.

The third statement follows from [7] where it is shown that an optimal Steiner forest can be computed in polynomial time if the underlying graph has treewidth at most 2. We show that every n -series-parallel graph is generalized series-parallel, and since these graphs have treewidth at most 2, the desired result follows. For a full proof of Theorem 10, see [34]. ◀

► **Remark.** The first two results of Theorem 10 can be generalized to nonnegative, non-decreasing and discrete-concave shareable edge cost functions. However, we do not know whether or not polynomial running time can be guaranteed.

We now demonstrate in Theorem 11 (for a proof see [34]) that the assumption of n -series-parallel graphs is in some sense well justified. Recall that a *generalized series-parallel graph* is created by a sequence of series, parallel, and/or add operations starting from a single edge, where an *add* operation adds a new vertex w and connects it to an existing vertex v by the edge $w - v$. The proofs of Theorem 10 and Theorem 11 show that the n -series-parallel graphs form a proper subclass of the generalized series-parallel graphs.

► **Theorem 11.** *For $n \geq 3$ players, there is a generalized series-parallel graph with fixed edge costs and no player-specific delays, so that the unique optimal Steiner forest is not enforceable. Therefore, a black-box reduction as for n -series-parallel graphs is impossible for generalized series-parallel graphs (even without player-specific delays).*

References

- 1 Nir Andelman, Michal Feldman, and Yishay Mansour. Strong price of anarchy. *Games Econom. Behav.*, 65(2):289–317, 2009.
- 2 Elliot Anshelevich and Bugra Caskurlu. Exact and approximate equilibria for optimal group network formation. *Theor. Comput. Sci.*, 412(39):5298–5314, 2011.
- 3 Elliot Anshelevich and Bugra Caskurlu. Price of stability in survivable network design. *Theory Comput. Syst.*, 49(1):98–138, 2011.
- 4 Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Tim Roughgarden, Éva Tardos, and Tom Wexler. The price of stability for network design with fair cost allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008. doi:10.1137/070680096.
- 5 Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4(1):77–109, 2008.
- 6 Guy Avni and Tami Tamir. Cost-sharing scheduling games on restricted unrelated machines. *Theor. Comput. Sci.*, 646:26–39, 2016.
- 7 Mohammadhossein Bateni, Mohammadtaghi Hajiaghayi, and Dániel Marx. Approximation schemes for steiner forest on planar graphs and graphs of bounded treewidth. *JACM*, 58(5):21:1–21:37, 2011.
- 8 Vittorio Bilò, Angelo Fanelli, Michele Flammini, and Luca Moscardelli. When ignorance helps: Graphical multicast cost sharing games. *Theoretical Computer Science*, 411(3):660–671, 2010.
- 9 Vittorio Bilò, Michele Flammini, and Luca Moscardelli. The price of stability for undirected broadcast network design with fair cost allocation is constant. In *Proc. 54th Symp. Foundations of Computer Science (FOCS)*, pages 638–647, 2013.
- 10 C. Bird. On cost allocation for a spanning tree: A game theoretic approach. *Networks*, 6:335–350, 1976.
- 11 Jaroslav Byrka and Karen Aardal. An optimal bifactor approximation algorithm for the metric uncapacitated facility location problem. *SIAM Journal on Computing*, 39(6):2212–2231, 2010.
- 12 Jaroslav Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. Steiner tree approximation via iterative randomized rounding. *J. ACM*, 60(1):6:1–6:33, 2013.
- 13 Ioannis Caragiannis, Vasilis Gkatzelis, and Cosimo Vinci. Coordination mechanisms, cost-sharing, and approximation algorithms for scheduling. In Nikhil R. Devanur and Pinyan Lu, editors, *Web and Internet Economics*, pages 74–87, 2017.
- 14 Jean Cardinal and Martin Hoefer. Non-cooperative facility location and covering games. *Theor. Comput. Sci.*, 411:1855–1876, March 2010.
- 15 Moses Charikar, Chandra Chekuri, To-Yat Cheung, Zuo Dai, Ashish Goel, and Sudipto Guha. Approximation algorithms for directed Steiner problems. *J. Algorithms*, 33(1):192–200, 1999.

- 16 Chandra Chekuri and Martin Pál. A recursive greedy algorithm for walks in directed graphs. In *Proc. 46th Symp. Foundations of Computer Science (FOCS)*, pages 245–253, 2005.
- 17 Ho-Lin Chen, Tim Roughgarden, and Gregory Valiant. Designing network protocols for good equilibria. *SIAM J. Comput.*, 39(5):1799–1832, 2010.
- 18 Xujin Chen, Zhuo Diao, and Xiaodong Hu. Network characterizations for excluding Braess’s paradox. *Theory Comput. Syst.*, 59(4):747–780, 2016.
- 19 Giorgos Christodoulou, Vasilis Gkatzelis, and Alkmini Sgouritsa. Cost-sharing methods for scheduling games under uncertainty. In *Proc. 18th ACM Conf. Economics and Computation (EC)*, pages 441–458, 2017.
- 20 Giorgos Christodoulou, Stefano Leonardi, and Alkmini Sgouritsa. Designing cost-sharing methods for bayesian games. In *Proc. 9th Intl. Symp. Algorithmic Game Theory (SAGT)*, pages 327–339, 2016.
- 21 Giorgos Christodoulou and Alkmini Sgouritsa. Designing networks with good equilibria under uncertainty. In *Proc. 27th Symp. Discrete Algorithms (SODA)*, pages 72–89, 2016.
- 22 Xiaotie Deng, Toshihide Ibaraki, and Hiroshi Nagamochi. Algorithmic aspects of the core of combinatorial optimization games. *Math. Oper. Res.*, 24(3):751–766, 1999.
- 23 Yann Disser, Andreas Feldmann, Max Klimm, and Matús Mihalák. Improving the H_k -bound on the price of stability in undirected shapley network design games. *Theoret. Comput. Sci.*, 562:557–564, 2015.
- 24 Michal Feldman and Tami Tamir. Conflicting congestion effects in resource allocation games. *Oper. Res.*, 60(3):529–540, 2012.
- 25 Amos Fiat, Haim Kaplan, Meital Levy, Svetlana Olonetzky, and Ronen Shabo. On the price of stability for designing undirected networks with fair cost allocations. In *Proc. 33rd Intl. Coll. Automata, Languages and Programming (ICALP)*, volume 1, pages 608–618, 2006.
- 26 Naveen Garg, Goran Konjevod, and R Ravi. A polylogarithmic approximation algorithm for the Group Steiner tree problem. *J. Algorithms*, 37:66–84, 2000.
- 27 Konstantinos Georgiou and Chaitanya Swamy. Black-box reductions for cost-sharing mechanism design. In *Proc. 23rd Symp. Discrete Algorithms (SODA)*, pages 896–913, 2012.
- 28 Michel X. Goemans and Martin Skutella. Cooperative facility location games. *J. Algorithms*, 50(2):194–214, 2004.
- 29 Daniel Granot and Gur Huberman. On minimum cost spanning tree games. *Math. Prog.*, 21:1–18, 1981.
- 30 Daniel Granot and Michael Maschler. Spanning network games. *Internat. J. Game Theory*, 27:467–500, 1998.
- 31 Sudipto Guha and Samir Khuller. Greedy strikes back: Improved facility location algorithms. *J. Algorithms*, 31:228–248, 1999.
- 32 Thomas Dueholm Hansen and Orestis Telelis. On pure and (approximate) strong equilibria of facility location games. In Christos Papadimitriou and Shuzhong Zhang, editors, *Internet and Network Economics*, pages 490–497, 2008.
- 33 Thomas Dueholm Hansen and Orestis Telelis. Improved bounds for facility location games with fair cost allocation. In *Proc. 3rd Intl. Conf. Combinatorial Optimization and Applications (COCOA)*, pages 174–185, 2009.
- 34 Tobias Harks, Martin Hoefer, Anja Huber, and Manuel Surek. Efficient black-box reductions for separable cost sharing, 2018. [arXiv:1802.10351](https://arxiv.org/abs/1802.10351) [cs.GT].
- 35 Tobias Harks, Anja Huber, and Manuel Surek. A characterization of undirected graphs admitting optimal cost shares. In Nikhil R. Devanur and Pinyan Lu, editors, *Web and Internet Economics*, pages 237–251, Cham, 2017. Springer International Publishing.
- 36 Tobias Harks and Britta Peis. Resource buying games. *Algorithmica*, 70(3):493–512, 2014.
- 37 Tobias Harks and Philipp von Falkenhausen. Optimal cost sharing for capacitated facility location games. *European Journal of Operational Research*, 239(1):187–198, 2014.

- 38 Martin Hoefer. Non-cooperative tree creation. *Algorithmica*, 53:104–131, 2009.
- 39 Martin Hoefer. Competitive cost sharing with economies of scale. *Algorithmica*, 60:743–765, 2011.
- 40 Martin Hoefer. Strategic cooperation in cost sharing games. *Internat. J. Game Theory*, 42(1):29–53, 2013.
- 41 Kamal Jain and Vijay Vazirani. Applications of approximation algorithms to cooperative games. In *Proc. 33rd Symp. Theory of Computing (STOC)*, pages 364–372, 2001.
- 42 Jochen Könemann, Stefano Leonardi, Guido Schäfer, and Stefan van Zwam. A group-strategyproof cost sharing mechanism for the steiner forest game. *SIAM J. Comput.*, 37(5):1319–1341, 2008.
- 43 Euiwoong Lee and Katrina Ligett. Improved bounds on the price of stability in network cost sharing games. In *Proc. 14th Conf. Electronic Commerce (EC)*, pages 607–620, 2013.
- 44 Jian Li. An $o(\log(n)/\log(\log(n)))$ upper bound on the price of stability for undirected shapley network design games. *Inform. Process. Lett.*, 109(15):876–878, 2009.
- 45 Shi Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Inf. Comput.*, 222:45–58, 2013.
- 46 Nimrod Megiddo. Cost allocation for Steiner trees. *Networks*, 8(1):1–6, 1978.
- 47 Hervé Moulin and Scott Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Econom. Theory*, 18(3):511–533, 2001.
- 48 Martin Pál and Éva Tardos. Group strategyproof mechanisms via primal-dual algorithms. In *FOCS*, pages 584–593, 2003.
- 49 Gabriel Robins and Alexander Zelikovsky. Tighter bounds for graph Steiner tree approximation. *SIAM J. Disc. Math.*, 19(1):122–134, 2005.
- 50 Robert Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory*, 2:65–67, 1973.
- 51 Vasilis Syrgkanis. The complexity of equilibria in cost sharing games. In A. Saberi, editor, *Proc. 6th Internat. Workshop on Internet and Network Econom.*, LNCS, pages 366–377, 2010.
- 52 Arie Tamir. On the core of network synthesis games. *Math. Prog.*, 50:123–135, 1991.
- 53 Philipp von Falkenhausen and Tobias Harks. Optimal cost sharing for resource selection games. *Math. Oper. Res.*, 38(1):184–208, 2013.