

Strategic Cooperation in Cost Sharing Games

Martin Hoefer

Received: date / Accepted: date

Abstract We examine strategic cost sharing games with so-called arbitrary sharing based on various combinatorial optimization problems. These games have recently been popular in computer science to study cost sharing in the context of the Internet. We concentrate on the existence and computational complexity of strong equilibria, in which no coalition can improve the cost of each of its members. Our main result reveals a connection to the core in coalitional cost sharing games studied in operations research. For set cover and facility location games this results in a tight characterization of the existence of strong equilibrium using the integrality gap of suitable linear programming formulations. Furthermore, it allows to derive all existing results for strong equilibria in network design cost sharing games with arbitrary sharing via a unified approach. In addition, we show that in general there is no efficiency loss, i.e., the strong price of anarchy is always 1. Finally, we indicate how the LP-approach is useful for the computation of near-optimal and near-stable approximate strong equilibria.

Keywords Cost Sharing · Strong Equilibrium · Integrality Gap · Combinatorial Optimization

1 Introduction

How can a set of self-interested actors share the cost of a joint investment in a fair and stable way? This fundamental question has motivated a large amount of research

Supported by DFG through UMIC Research Centre at RWTH Aachen University and grant Ho 3831/3-1. An extended abstract of this work appeared in the Proceedings of the 6th Workshop on Internet & Network Economics (Hoefer, 2010).

RWTH Aachen University
Lehrstuhl Informatik I
52074 Aachen, Germany
Tel.: +49-241-8021104
Fax: 49-241-8022216
E-mail: mhoefer@cs.rwth-aachen.de

in economics in the last decades (Young, 1994). More recently, this question has been studied in computer science to understand the development of the Internet and questions arising in e-commerce (Jain and Mahdian, 2007). A classic framework to study cost sharing problems without centralized control are *cost sharing games*, in which cost can be specified as an abstract parameter for each player and/or each coalition. Relevant to real-world optimization problems are cost sharing games, where the cost is tied to the investment into specific resources. Such games based on combinatorial optimization problems have a long tradition in economics and operations research. They are usually formulated as coalitional game, i.e., there is a set of players, and each coalition of players has an associated cost value. This value comes from an optimal solution to an underlying optimization problem for the coalition. For example, consider a multicast network design game, in which players in a network strive to receive a message by establishing a connection to a common source vertex s . A simple model for this scenario is a *minimum spanning tree game (MST game)*, in which each vertex $v \neq s$ is a player, and the edges have costs. The cost of a coalition C in this game is the cost of the cheapest network spanning all players in C and s , i.e., the minimum spanning tree for the set $C \cup s$. In the literature many interesting and important coalitional cost sharing games have been studied, e.g., based on problems like MST (Bird, 1976) and Steiner tree (Megiddo, 1978; Granot and Huberman, 1981; Tamir, 1991; Granot and Maschler, 1998), covering and packing problems (Deng et al., 1999), facility location (Tamir, 1993; Goemans and Skutella, 2004), or TSP (Faigle et al., 1998). Of central importance in these games is the existence and computation of a stable and fair cost sharing among the players. Coalitional cost sharing games are usually *transferable utility (TU) games*, i.e., the cost can be shared arbitrarily between the players. This allows for the largest level of generality for possible interactions in the bargaining and coalition formation process. The foremost concept of stability and fairness in TU cost sharing games is the *core*. The core is a set of imputations, i.e., of distributions of the cost for the complete player set to the players. To be in the core an imputation has to fulfill the additional property that no coalition of players in sum pays more than its associated cost value. Results about the non-emptiness of the core and characterizations of core solutions have been obtained for many of the games mentioned above.

A problem with coalitional cost sharing games is that cost shares represent a strong abstraction from the underlying optimization problem. Players are assumed to contribute on a global level, and the game does not take into account who pays how much for which resource. In particular, our simple MST example above is appropriate when we think of a set of players signing up for message reception with a service provider. The provider then decides which links are established and how the network is built to serve all users that signed up. Finally, he simply charges the players a total cost and they have to come up with a way to pay the bill. Here players are not in charge of the decision about which edges to create and spend their money on. However, when studying the incentives in large unregulated settings, such as, e.g., in the creation of the Internet, there is a need to understand cost sharing on a more detailed and, in particular, a strategic level. This prompted computer scientists to study strategic cost sharing scenarios. On the one hand, there are a number of recent works on *designing* strategic cost sharing games to obtain favorable Nash equilibrium proper-

ties (Chen et al., 2010). In these games the underlying assumption is that a central authority designs and maintains the solution and dictates cost shares for each player. This is actually quite close to the underlying assumptions in, e.g., coalitional games with non-transferable utility or cost sharing mechanisms, which have received a lot of attention (Immorlica et al., 2008; Jain and Vazirani, 2001; Könemann et al., 2008; Pál and Tardos, 2003). Designing cost shares, e.g., using Shapley value cost sharing (Albers, 2009; Anshelevich et al., 2008a; Epstein et al., 2009; Leonardi and Sankowski, 2007), can yield favorable properties concerning existence and cost of Nash equilibria. In contrast, such a model is unsuitable when there is very little control over players and their bargaining options. A model that allows for general cost sharing between players is sometimes referred to as *arbitrary* cost sharing, and it has been studied widely (Anshelevich et al., 2008b; Anshelevich and Caskurlu, 2011a,b; Cardinal and Hoefer, 2010; Epstein et al., 2009; Hoefer, 2009, 2011; Hoefer and Krysta, 2005). In these cost sharing games the strategy of a player is a payment function that specifies his exact contribution to the cost of each resource. The outcomes of such strategic cost sharing games based on combinatorial optimization problems will be the subject of this paper.

When studying the outcomes of the interaction of rational agents in strategic games, we need to discuss the appropriate solution concept. The most prominent stability concept in strategic games is the Nash equilibrium (NE). While a mixed NE in finite games always exists, a drawback is that it is only resilient to unilateral deviations. In many reasonable scenarios agents might be able to coordinate their actions, and under these circumstances a NE is not a reasonable solution concept. To address this issue we consider the *strong equilibrium* (SE) in this paper. A strong equilibrium (Aumann, 1959) is a state in which no coalition (of arbitrary size) has a deviation that lowers the cost of *every* member of the coalition. This resilience to coalitional deviations is highly attractive. On the downside, strong equilibria might not exist in a game. This may be the reason they have not received an equivalent amount of research interest despite their attractive properties. We partly circumvent this problem by studying approximate versions of the strong equilibrium, which is guaranteed to exist. However, our treatment of these aspects is brief and mostly left for future work.

Our main interest is to characterize the existence, social cost, and computational complexity of SE in strategic cost sharing games based on combinatorial optimization problems. An initial insight in Section 2 reveals that the concept of SE in strategic cost sharing games is equivalent to seemingly stronger notions of super-strong or sum-strong equilibria. Additionally, a SE in a strategic game can always be turned into a core imputation of the corresponding coalitional game defined on the same instance of the optimization problem. Hence, a SE represents a *strategic refinement* of a core solution, and existence of a SE implies non-emptiness of the core. It also implies that the strong price of anarchy (Andelman et al., 2009) is 1, i.e., in every SE a solution is bought that is a social optimum to the underlying optimization problem.

In Section 3 we consider a variety of games based on vertex and set cover and various facility location problems. For these games we show an equivalence result of core and SE. In particular, whenever the core in the coalitional game is nonempty, there is a SE for the strategic game. Our main proof technique relies on a connection

via linear programming to Owen’s linear production model (Owen, 1975), which is one of the most common ways to show non-emptiness of the core in combinatorial optimization games (Deng et al., 1999). Using this machinery we are able to tightly characterize the existence and cost of SE in all these games. Our results extend to special classes of network design problems. This includes, e.g., MST and classes of Steiner Network Design Games (Anshelevich et al., 2008b; Hoefer, 2009). As a byproduct, our general approach yields simple proofs for all known results for SE in strategic cost sharing games with arbitrary sharing, which were previously shown by Epstein et al. (2009) via complicated combinatorial arguments.

The equivalence between SE and core solutions is an interesting and notable fact. However, in other cases such as NTU games and appropriate extensions to strategic games a similar equivalence is obvious. Thus, it may be more surprising to observe that the relation between SE and core solutions in cost sharing games can be quite complicated. In particular, in Section 4 we explore equivalence without relying on linear programming. While in some cases like Terminal Backup Games (Anshelevich and Caskurlu, 2011b) we can resort to combinatorial arguments, in other interesting games our results are mostly negative. In particular, we show that in Steiner Tree connection games or network cutting games equivalence does not hold, i.e., the core might be non-empty but a SE is absent. A similar result is established in Section 5.1 even for simple vertex cover games when we allow resources to be purchased fractionally or in multiple units. Characterizing SE in these games remains as an intriguing open problem. We observe in Section 5.2 that linear programming can be used to obtain approximate SE in vertex and set cover, as well as facility location games. Finally, we conclude in Section 6 with some interesting questions for further research.

Our main conceptual contribution is to reveal a non-trivial and close relation between coalitional and strategic games defined on the same instance of the optimization problem. The strategic game can be seen as a strategic variant of the coalitional game. In addition, in many games SE can act as a strategic refinement of rather coarse core solutions. We believe that this inherent connection should stimulate further research on (strategic) cost sharing with rational agents.

1.1 Preliminaries

We consider classes of cost sharing games based on combinatorial optimization problems. In each of these games there is a set R of resources. Resource $r \in R$ can be *bought* if the associated cost $c(r) \geq 0$ is paid for. For $R' \subseteq R$ let $c(R') = \sum_{r \in R'} c(r)$. We assume that there is set of players K . Each player $i \in K$ strives to satisfy a certain constraint on the bought resources. For example, in the case of the *set cover problem* the player set is the element set $K = E$. The resources are sets $R = \mathcal{S} \subseteq 2^E$ over E . The constraint of player e states that there must be at least one bought set S with $e \in S$. In a similar way we can base our construction on various cost minimization problems like facility location or network design. We will describe them in more detail in the corresponding sections. However, a common assumption in our problems is a free disposal property, i.e., if for a set of bought resources all player constraints

are satisfied, then a superset of bought resources can never make a player constraint become violated.

For a given set of players, resources, and constraints we define two games - a *coalitional* and a *strategic* cost sharing game. The *coalitional game* $\Delta = (K, c)$ is given by the set of players K and a cost function $c : 2^K \rightarrow \mathbb{R}_0^+$ that specifies a cost value for every subset of players. For a coalition $C \subseteq K$, the cost is $c(C) = c(R(C)^*) = \sum_{r \in R(C)^*} c(r)$ for an *optimum solution* $R(C)^* \subseteq R$ for C . In particular, $R(C)^*$ is a minimum cost set of resources that must be bought to satisfy all constraints of players in C . For example, in a set cover game $R(C)^*$ is the minimum cost set cover for the elements in C . We denote the special case $R^* = R(K)^*$ as the *social optimum*.

The goal in a coalitional game is to find a cost sharing of $c(K)$ for the so-called grand coalition K . A vector of cost shares $\gamma_1, \dots, \gamma_k$ is called an *imputation* if $\sum_{i \in K} \gamma_i = c(K)$. The game Δ is a *transferable utility (TU)* game, i.e., we are free to choose $0 \leq \gamma_i \leq c(K)$. The central concept of stability and fairness in coalitional games is the *core*. The core is the set of imputations γ , for which $c(C) \geq \sum_{i \in C} \gamma_i$. Intuitively, when sharing the cost according to a member of the core, no subset of players has an incentive to deviate from the grand coalition and make a separate investment - depending on the underlying optimization problem, e.g., purchase different sets or construct an independent network.

The *strategic game* $\Gamma = (K, (S_i)_{i \in K}, (c_i)_{i \in K})$ is specified by strategies and individual cost for each player. The *strategy space* S_i of player $i \in K$ consists of all functions $s_i : R \rightarrow \mathbb{R}_0^+$. Strategy s_i allows him to specify for each resource $r \in R$ how much he is willing to contribute to r . A resource r is *bought* if $\sum_{i \in K} s_i(r) \geq c(r)$. A vector of strategies s is a *state* of the game. For a state s we define $|s_i| = \sum_{r \in R} s_i(r)$ and the *individual cost* of player i as $c_i(s) = |s_i|$ if the bought resources satisfy his constraint. Otherwise, $c_i(s) = \infty$ or a different value that is prohibitively large. Finally, the *social cost* of s is $c(s) = \sum_{i \in K} c_i(s)$.

The foremost concept of stability in strategic games is the Nash equilibrium, a state in which no player unilaterally has an incentive to deviate. In this paper, however, we consider coalitional incentives and thus resort to a strengthened version called strong equilibrium (Aumann, 1959). A state s has a *violating coalition* $C \subseteq K$ if there are strategies $s'_C = (s'_i)_{i \in C}$ such that $c_i(s'_C, s_{-C}) < c_i(s)$ for each $i \in C$. A violating coalition has a deviation, in which all players in C strictly pay less. A *strong equilibrium* is a state s that has no violating coalition. Note that in a SE a set of resources is bought such that all player constraints are satisfied. Each resource r is either paid for exactly or not contributed to at all. Thus, a SE represents a cost sharing of some feasible solution R for the grand coalition, such that $c(s) = c(R)$.

In addition, we briefly consider the concept of an (α, β) -approximate strong equilibrium (denoted (α, β) -SE). A state s is a (α, β) -SE if for every coalition $C \subseteq K$ and strategies s'_C there is at least one $i \in C$ such that $c_i(s) \leq \alpha \cdot c_i(s'_C, s_{-C})$, and if $c(s) \leq \beta \cdot c(K)$. In such a state no coalition can reduce the cost of every member by strictly more than a factor of α , and the cost of the bought solution represents a β -approximation to $c(K)$.

2 Strong Equilibria and the Core

Consider a given set of resources R with costs $c(r)$ and a set of players K with constraints. Our first insight reveals that several coalitional equilibrium concepts coincide in strategic games Γ . In particular, we consider *super-strong* and *sum-strong equilibria* defined as follows. A state s has a *weakly violating coalition* $C \subseteq K$ if there are strategies $s'_C = (s'_i)_{i \in C}$ such that $c_i(s'_{C, s-C}) \leq c_i(s)$ for each $i \in C$ and $c_{i'}(s'_{C, s-C}) < c_{i'}(s)$ for at least one $i' \in C$. A state s has a *sum violating coalition* $C \subseteq K$ if there are strategies $s'_C = (s'_i)_{i \in C}$ such that $\sum_{i \in C} c_i(s'_{C, s-C}) < \sum_{i \in C} c_i(s)$. A super-strong (sum-strong) equilibrium is a state s that has no weakly (sum) violating coalition. Note that every violating coalition is also weakly violating, and every weakly violating coalition is also sum violating. Hence, every sum-strong equilibrium is super-strong, and every super-strong equilibrium is strong. We note this simple fact because we actually show the absence of sum violating coalitions in our proofs below. In general strategic games it is easy to see that the inclusions are strict, i.e., a strong equilibrium might not be sum-strong. In our strategic cost sharing games Γ , however, every strong equilibrium is also sum-strong.

Proposition 1 *Every strong equilibrium in a strategic game Γ is a sum-strong equilibrium.*

Proof Suppose an arbitrary state of the strategic game Γ has a sum violating coalition that can achieve a strict improvement in the sum of player costs. We will show that in this case it also has a violating coalition that can obtain a strict improvement for every player of the coalition. The proposition then shows that strong equilibria without violating coalitions are also sum-strong.

Consider a strategic game Γ and an arbitrary state s that has a sum violating coalition C with a deviation s'_C , i.e., $\sum_{k \in C} c_k(s'_{C, s-C}) < \sum_{k \in C} c_k(s)$. The cost of a player k is either his total payment $|s_k|$ or ∞ . In $(s'_{C, s-C})$ none of the players in C can have cost ∞ , because then the sum of costs cannot represent a strict improvement over that in s . Hence, for all $k \in C$ we must have finite cost $c_k(s'_{C, s-C}) = |s'_k|$ in $(s'_{C, s-C})$, and hence, the set of bought resources in $(s'_{C, s-C})$ satisfies the constraint of every player in C .

First, suppose that in s there is a player $k \in K$ with cost $c_k(s) = \infty$. This means his constraint is not satisfied by the bought resources in s . k has a unilateral deviation of purchasing all resources in R^* by himself. This yields finite cost for k . Hence, $\{k\}$ is both, a (singleton) violating coalition and a sum violating coalition.

Second, suppose that in s for all players $k \in K$ we have finite cost $c_k(s) = |s_k|$. For the sum violating coalition C we consider the subset of players C_0 with $|s_k| = 0$. A player in C_0 cannot achieve a strict improvement, so he cannot be part of a violating coalition in which *every player strictly* improves. Instead, we prove that $C_1 = C - C_0$ is a violating coalition. Note that C_1 must be non-empty, because all players in K have finite cost in s . In this case, $C_1 = \emptyset$ would imply that $c_k(s) = 0$ for all $k \in C$, which is impossible to improve and contradicts that C is a sum violating coalition.

We now construct a deviation strictly improving the cost for every player in C_1 as follows. Let $s'(r) = \sum_{j \in C} s'_j(r)$ be the total contribution by players in C to resource $r \in R$ in the deviation s' . For each player $k \in C$ we define a strategy $s''_k(r) =$

$(|s_k|/\sum_{j \in C} |s_j|) \cdot s'(r)$. Thus, in sum players in C contribute the same to r in s' and s'' , however, in s'' each player k pays a share at each resource that corresponds to the fraction of total cost contributed by k in s . Note that this yields $|s''_k| = 0$ if and only if $|s_k| = 0$, as in this case the fraction of player k is 0. Thus, we can let players in C_0 stick to their original strategy and concentrate on players in C_1 for the switch to $s''_{C_1} = (s''_i)_{i \in C_1}$. For every player $k \in C_1$ we have

$$|s''_k| = \sum_{r \in R} \frac{|s_k|}{\sum_{j \in C} |s_j|} \cdot s'(r) = |s_k| \cdot \frac{\sum_{j \in C} |s'_j|}{\sum_{j \in C} |s_j|} = |s_k| \cdot \frac{\sum_{j \in C} c_j(s'_C, s_{-C})}{\sum_{j \in C} c_j(s)} < |s_k| ,$$

using the assumption that (s'_C, s_{-C}) strictly improves the sum of costs for players in C . Note that for each resource r we have the same total contribution from s'_C and s''_{C_1} . Thus, in (s''_{C_1}, s_{-C_1}) the same resources get bought as in (s'_C, s_{-C}) and, as observed above, they satisfy the constraint for every player $k \in C$. Hence, for every $k \in C_1$ we have $c_k(s''_{C_1}, s_{-C_1}) = |s''_k| < |s_k| = c_k(s)$. This proves that C_1 is a violating coalition. Thus, for every sum violating coalition C there is a violating coalition C_1 . This proves the proposition. \square

We note on the side that for every state s , coalition C , deviation $s'_C = (s'_i)_{i \in C}$, and finite $\alpha \geq 1$ with $\sum_{i \in C} c_i(s'_C, s_{-C}) = \alpha \cdot \sum_{i \in C} c_i(s)$ we can find in a similar way C' and $s''_{C'} = (s''_i)_{i \in C'}$ with $c_i(s''_{C'}, s''_{-C'}) = \alpha c_i(s)$ for every $i \in C'$. Thus, the equivalence of strong and sum-strong equilibria holds also for approximate versions of the concepts, in which players must improve their costs by a factor of strictly more than α .

We continue to show a general connection between core imputations for the coalitional game Δ and SE of the strategic game Γ . We first observe that in a SE players always share the cost of a social optimum R^* .

Proposition 2 *In every strong equilibrium of a strategic game Γ the players share the cost of a social optimum. The strong price of anarchy is 1.*

Proof Consider a SE s and the set R' of bought resources. Assume for contradiction $c(R') > c(R^*)$. If all players with $|s_k| > 0$ jointly deviate to purchase R^* , each player k must pay only a fraction of $c(R^*)/c(R') < 1$ of $|s_k|$. Formally, define $s'_k(r) = c(r) \cdot \frac{|s_k|}{c(R')}$. If all contributing players jointly deviate to s' , this obviously strictly decreases the payment of *all* players. Hence, if $c(R') > c(R^*)$, then K is a violating coalition for s , a contradiction. \square

Proposition 3 *If the strategic game Γ has a strong equilibrium, then the coalitional game Δ has a core solution.*

Proof Consider a SE s of Γ , which by Proposition 2 is a cost sharing of R^* , and a coalition C . The coalition has the possibility to deviate and contribute just to buy $R(C)^*$. In this case it has to share for each resource $r \in R(C)^*$ at most the remaining cost on top of the contribution of players in $K \setminus C$, i.e., $c_C(r) = c(r) - \sum_{k \in K \setminus C} s_k(r)$. If $c_C(R(C)^*) < \sum_{k \in C} |s_k|$, the coalition can deviate to $s'_k(r) = c_C(r) \cdot \frac{|s_k|}{\sum_{j \in C} |s_j|}$ for every $k \in C$ and every $r \in R(C)^*$, which would represent an improvement for every player in C . However, as s is a SE, C must not be violating, and so $c_C(R(C)^*) \geq \sum_{k \in C} |s_k|$.

Trivially, $c(r) \geq c_C(r)$, and so $c(R(C)^*) \geq \sum_{k \in C} |s_k|$. Thus, γ with $\gamma_k = |s_k|$ is in the core of Δ . \square

Intuitively, this shows that the core is less stringent as we assume in a deviation of a coalition C that players outside C “stop contributing”. More formally, non-emptiness of the core is a necessary condition for existence of a SE. In the following we consider various classes of games, in which it is also sufficient. In these cases the SE is a strategic refinement of the core, as it allows to specify a strategic allocation of payments to resources.

3 Strong Equilibria using Linear Programming

3.1 Vertex and Set Cover Games

In a variety of fundamental games non-emptiness of the core and existence of SE are equivalent. We can relate SE existence to the core via linear programming duality. For simplicity we outline the general argument in the setting of set cover games. In a set cover game, we are given a set of players as elements E and a set system $\mathcal{S} \subseteq 2^E$, where each $S \in \mathcal{S}$ has a cost $c(S) \geq 0$. The constraint of player e is that at least one set S with $e \in S$ must be bought.

Theorem 1 *If a set cover game Δ has a non-empty core, then the strategic game Γ has a strong equilibrium.*

Proof We consider the integer programming formulation of set cover. In particular, we consider the following linear relaxation, which employs $x_S \geq 0$ instead of $x_S \in \{0, 1\}$ and thus allows sets to be included fractionally in the solution.

$$\begin{aligned} \text{Min} \quad & \sum_{S \in \mathcal{S}} x_S c(S) \\ \text{subject to} \quad & \sum_{S: e \in S} x_S \geq 1 \quad \forall e \in E \\ & x_S \geq 0 \quad \forall S \in \mathcal{S}. \end{aligned}$$

We also consider the corresponding LP dual.

$$\begin{aligned} \text{Max} \quad & \sum_{e \in E} \gamma_e \\ \text{subject to} \quad & \sum_{e \in S} \gamma_e \leq c(S) \quad \forall S \in \mathcal{S} \\ & \gamma_e \geq 0 \quad \forall e \in E. \end{aligned}$$

It has been shown by Deng et al. (1999) that the core of Δ is non-empty if and only if the integrality gap of the underlying set cover problem is 1, i.e., if the LP has an integral optimal solution. With Proposition 3 this is a prerequisite for existence of a SE in Γ . We strengthen this result by showing that core solutions can also be turned into an allocation of payments to resources for a SE in Γ . Thus, an integral optimum is also sufficient.

For the above programs consider the optimum primal solution x^* and the optimum dual solution γ^* , where x^* is integral and defines a feasible cover. Both x^* for the primal and γ^* for the dual yield the same objective value. Now assign each player e to pay $s_e(S) = \gamma_e^* x_S^*$ if $e \in S$ and $s_e(S) = 0$ otherwise. The theorem follows if every set in the cover is purchased exactly and no coalition C can reduce their total payments $\sum_{e \in C} |s_e|$. The first condition is clearly necessary for a SE, the second one implies that no coalition can be sum violating (and thus violating). We first show that the sets are exactly paid for. If $x_S^* > 0$, then due to complementary slackness the inequality $\sum_{e \in S} \gamma_e^* \leq c(S)$ is tight, hence by this assignment all the purchased sets get exactly paid for.

We now show that no coalition can reduce the total payments. The main idea of this part is to use duality arguments for a cost reduction of resources. In particular, for an optimum x^* , the objective function can be represented by a linear combination of tight constraints. The multipliers are the optimal dual variables γ^* . Due to complementary slackness, we can replace each $c(S)$ of a bought set S in the objective of the primal by $c(S) = \sum_{e \in S} \gamma_e^*$. For every coalition, this additive structure allows to reduce the costs and drop the shares bought by other players outside the coalition. In this way, we can show optimality of x^* under the remaining costs for every coalition and contradict that a coalition is sum violating or violating.

In particular, suppose for contradiction there is a coalition C that is sum violating, i.e., it has a deviation to strictly reduce their total payments. To find a deviation for the coalition that strictly improves their total cost, we formulate the optimization problem of finding a minimum cost cover for coalition C given the contributions of players $e \notin C$. The players in C can use the contributions by players in $K - C$, and thus for C the cost of a set S becomes

$$c_C(S) = c(S) - \sum_{e \notin C, e \in S} \gamma_e^* x_S^* .$$

Finding a minimum cost cover for coalition C with these adjusted costs can be formulated by the following reduced primal LP

$$\begin{aligned} \text{Min} \quad & \sum_{S \in \mathcal{S}} x_S c_C(S) \\ \text{subject to} \quad & \sum_{S: e \in S} x_S \geq 1 \quad \forall e \in C \\ & x_S \geq 0 \quad \forall S \in \mathcal{S}. \end{aligned}$$

Note that for this reduced LP the solution x^* is obviously still feasible, because we only removed all constraints for elements $e \notin C$.

The dual of this program is

$$\begin{aligned} \text{Max} \quad & \sum_{e \in C} \gamma_e \\ \text{subject to} \quad & \sum_{e \in S} \gamma_e \leq c_C(S) \quad \forall S \in \mathcal{S} \\ & \gamma_e \geq 0 \quad \forall e \in C. \end{aligned}$$

Note that the constraints of this program read

$$\sum_{e \in C, e \in S} \gamma_e + \sum_{e \notin C, e \in S} \gamma_e^* x_S^* \leq c(S) .$$

Setting $\gamma_e = \gamma_e^*$ for all $e \in C$ yields a feasible solution to the LP-dual of the reduced problem, because $x_S^* \leq 1$ and γ^* was feasible for the original dual.

We now observe that the objective function value of x^* and γ^* for the reduced problems is the same by using the integrality of x^* and the additive decomposition of costs resulting from complementary slackness. In particular, x^* has a value of

$$\begin{aligned} \sum_{S \in \mathcal{S}} x_S^* \left(c(S) - \sum_{e \notin C, e \in S} \gamma_e^* x_S^* \right) &= \sum_{S: x_S^* = 1} c(S) - \sum_{e \notin C, e \in S} \gamma_e^* \\ &= \sum_{S: x_S^* = 1} \sum_{e \in C, e \in S} \gamma_e^* \\ &= \sum_{e \in C} \gamma_e^* . \end{aligned}$$

The first equality follows because x^* is binary. The second equality follows because $c(S) = \sum_{e \in S} \gamma_e^*$ when $x_S^* = 1$ due to complementary slackness for the original LPs. Finally, the third equality is again due to complementary slackness, because $\gamma_e^* = 0$ whenever $\sum_{S: e \in S} x_S^* > 1$.

Hence, x^* and γ^* are both feasible for the reduced primal and dual programs and they yield the same value of the objective function. By strong duality both x^* and γ^* must be optimal solutions to the reduced primal and dual problems. In particular, x^* being an optimal solution to the reduced problem implies that the coalition C achieves minimum total payments by paying the remaining cost of the sets bought in x^* . Hence, C cannot be sum violating and not violating, a contradiction. This proves that s is a SE. \square

For the special case of vertex cover games we can use results from Deng et al. (1999) to efficiently compute SE. In particular, a game allows a core solution (and thus a SE) if and only if a maximum matching in the graph has the same size as the minimum vertex cover. This condition can be checked in polynomial time by computing corresponding vertex covers and matchings (Deng et al., 1999, Theorem 7 and Corollary 7). Hence, we can check in polynomial time whether a SE exists. If it exists, we can use the computed vertex cover as primal solution for our LP and compute cost shares for a strong equilibrium with the corresponding dual solution.

Corollary 1 *In a vertex cover game we can decide in polynomial time if a strong equilibrium exists. If it exists, we can compute a strong equilibrium in polynomial time.*

In addition, we can check in polynomial time whether a given strategy profile is a SE.

Corollary 2 *Given a state s for a vertex cover game Γ we can verify in polynomial time if it is a strong equilibrium.*

If the strategy profile is a SE, it must exactly pay for a vertex cover of the problem. This yields a primal solution for the LP. In addition, the accumulated cost shares of players must yield a corresponding dual solution. Finally, both primal and dual solutions must generate the same value of the objective function. This is a sufficient and necessary condition for being a SE, which can be checked in polynomial time.

Another interesting case are edge cover games. Here players are the vertices of a graph and resources are the edges. Each vertex wants to ensure that at least one incident edge is bought. Using the characterization of the non-emptiness of the core in Deng et al. (1999, Theorem 8 and Corollary 8) we can obtain similar results for this game as well.

Corollary 3 *In an edge cover game Γ we can decide in polynomial time if a strong equilibrium exists. If it exists, we can compute a strong equilibrium in polynomial time. Given a state s for an edge cover game Γ we can verify in polynomial time if it is a strong equilibrium.*

3.2 Facility Location Games

Another class of games that can be handled via similar arguments are facility location games. We outline the arguments on the simple class of *uncapacitated facility location games* (UFL games) and show below how to extend this approach to a more general class of games considered by Goemans and Skutella (2004) and Cardinal and Hoefer (2010). In a *UFL problem* there is a set T of terminals and a set F of facilities. We set $n_t = |T|$ and $n_f = |F|$. Each facility $f \in F$ has an opening cost $c(f) \geq 0$, for each terminal $t \in T$ and each facility $f \in F$ there is a connection cost $c(t, f) \geq 0$. The goal is to open a subset of facilities and buy a set of connections of minimum total cost, such that each terminal is connected to an opened facility. In the *UFL game* each player owns a terminal, i.e., $K = T$. The constraint of player t is satisfied if there is a bought connection (t, f) to some opened facility f . We can formalize the UFL problem by an integer program as follows:

$$\begin{aligned} \text{Min} \quad & \sum_{f \in F} c(f)y_f + \sum_{t \in T} c(t, f)x_{tf} \\ \text{subject to} \quad & \sum_{f \in F} x_{tf} \geq 1 \quad \text{for all } t \in T \\ & y_f - x_{tf} \geq 0 \quad \text{for all } t \in T, f \in F \\ & y_f, x_{tf} \in \{0, 1\} \quad \text{for all } t \in T, f \in F, \end{aligned}$$

Theorem 2 *If a UFL game Δ has a non-empty core, then the strategic game Γ has a strong equilibrium.*

Proof We again use the linear relaxation, which can be obtained by replacing $y_f, x_{tf} \in \{0, 1\}$ by $y_f, x_{tf} \geq 0$. Then the dual can be given by

$$\begin{aligned} \text{Max} \quad & \sum_{t \in T} \gamma_t \\ \text{subject to} \quad & \gamma_t - \delta_{tf} \leq c(t, f) \quad \text{for all } t \in T, f \in F \\ & \sum_{t \in T} \delta_{tf} \leq c(f) \quad \text{for all } f \in F. \end{aligned}$$

It has been shown in Goemans and Skutella (2004) that the core of Δ is non-empty if and only if the integrality gap of this LP is 1. We can now argue similarly as before. An integral optimum solution (x^*, y^*) to the LP-relaxation represents a partition of the terminal set T into a collection of stars, one for each facility f . The constraints

corresponding to these sets hold with tightness, and we can assign each player t to pay for her terminal the amount $s_t(t, f) = (\gamma_t^* - \delta_{tf}^*)x_{tf}^*$ as connection cost to f , in which (γ^*, δ^*) is the optimum solution to the dual. For the opening costs $s_t(f) = \delta_{tf}^*y_f^*$. In total this pays exactly for all costs of the solution by duality.

Suppose there is a violating coalition C . We again remove players in $K - C$ and reduce the costs of connections and facilities by the respective contributions. In order to represent a violating coalition, the players in C must be able to deviate and reduce their total sum of payments. However, the solution (x^*, y^*) has the same value for the reduced LP of coalition C as (γ^*, δ^*) for the dual of the reduced LP. By duality both solutions remain optimal. Thus, coalition C is purchasing an optimal solution against the payments of players in $K - C$ and has no possibility to reduce the total payments. This is a contradiction to C being a violating coalition. \square

This result can be used to characterize computational properties of SE. In particular, we can decide in polynomial time if a given strategy profile for Γ is a SE. We first check if the payments of players are made only to their own connection and opening costs. Then we accumulate contributions to cost shares and check if this yields a core solution - i.e., if the primal solution (given by the purchased solution to the facility location problem) and the dual solution (given by the cost shares) correspond to each other and yield the same optimal value for primal and dual LPs.

Corollary 4 *Given a strategy vector for a UFL game Γ we can verify in polynomial time if it is a strong equilibrium.*

As verification is in P, the problem of computing a strong equilibrium is in NP. In fact, Goemans and Skutella (2004) show for a class of UFL games that deciding the existence of a core solution is NP-complete. As existence of SE and core solutions is equivalent, this yields the following result.

Corollary 5 *It is NP-complete to decide if a given UFL game Γ has a strong equilibrium.*

This main arguments from the proofs above can be extended to the class of connection-restricted facility location games (CRFL games), in which access to a facility f can be obtained only by certain allowed coalitions $\mathcal{A}_f \subseteq 2^T$. We consider the special case of *closed* games (CCRFL games), in which the set system \mathcal{A}_f of allowed coalitions is downward closed, i.e., each subset of an allowed coalition is also an allowed coalition. This simplifies the specific allocation of the cost shares to connections and facilities. While the closed property is a restriction, we note that many variants of facility location arising in practice fall into this class of games, e.g., problems with capacity or incompatibility constraints. We believe that equivalence between core and SE also holds for CRFL games in full generality, but a proof of this statement remains as an open problem. For formal discussion and the proof of the following theorem see the Appendix.

Theorem 3 *If a CCRFL game Δ has a non-empty core, then the strategic game Γ has a strong equilibrium. Given a strategy vector for a CCRFL game Γ we can verify in polynomial time if it is a strong equilibrium.*

3.3 Connection Games

In this section we use a linear program to formulate network design games in directed and undirected graphs. Perhaps the most frequently studied variant is a *connection game* originally formulated by Anshelevich et al. (2008b). In this game there is a graph $G = (V, E)$, resources are the edges, and each edge has a non-negative cost $c(e) \geq 0$. There is a set of players K , and each player k has a source-sink pair (s_k, t_k) . A player is satisfied if there is a path of bought edges connecting his source and sink. This is a game based on the Steiner network problem in graphs (Goemans and Williamson, 1995). In a variant based on Steiner Tree called the *single-source* game, every player has the same source s . Here we characterize existence of SE based on a Flow-LP previously studied (Tamir, 1991; Wong, 1984).

Theorem 4 *If the Flow-LP has an integral optimum solution, then the strategic connection game Γ has a strong equilibrium.*

Proof We formulate the mixed integer program (MIP) for the problem in directed graphs. It is simple to adjust it to undirected graphs, where we use only one variable y_{ij} for each (undirected) edge $e = (i, j) \in E$.

$$\begin{aligned}
& \text{Min} \quad \sum_{(i,j) \in E} c_{ij} y_{ij} \\
& \text{s.t.} \quad \sum_{\{j \mid (i,j) \in E\}} f_{ij}^k - \sum_{\{j \mid (j,i) \in E\}} f_{ji}^k \geq 1 \quad \text{for } i = s_k \\
& \quad \quad \sum_{\{j \mid (i,j) \in E\}} f_{ij}^k - \sum_{\{j \mid (j,i) \in E\}} f_{ji}^k \geq 0 \quad \text{for } i \neq s_k, t_k \\
& \quad \quad y_{ij} - f_{ij}^k \geq 0 \quad \quad \text{for } (i, j) \in E, k \in K \\
& \quad \quad f_{ij}^k \geq 0, y_{ij} \in \{0, 1\} \quad \text{for } (i, j) \in E, k \in K
\end{aligned}$$

In this MIP we optimize for each player k a flow, which is required to have value 1 by the constraints at the source, and which can only exit through the sink. The individual flows are coordinated by capacity constraints $y_{ij} - f_{ij}^k \geq 0$. Each edge that is used by at least one player fractionally has to be fully paid for in the objective function. We can relax this program by using $y_{ij} \geq 0$. Then the dual can be formulated using variables δ_i^k for the flow conservation constraints and γ_{ij}^k for the coordination constraints. Intuitively, the values δ_i^k introduce a node potential of contributions, and γ_{ij}^k can be seen as contributions towards the edges that are bought.

It has been observed by Tamir (1991) that this program is within Owens linear production model. Hence, if the integrality gap is 1, the optimal dual solution yields a core solution. Using similar arguments as before, we can also show that in this case a SE exists. In particular, each player pays $s_k(i, j) = y_{ij}^* (\gamma_{ij}^{k*})$ towards edge (i, j) . For a coalition C we can again reduce costs of edges by removing players of $K - C$. Due to the additive structure of the LP, the primal and dual optimal solutions remain optimal for the reduced LPs. This means no coalition can reduce total payments, and no coalition can be violating. This proves the theorem. \square

This insight allows us to derive one of the main results shown by Epstein et al. (2009) in a simple and compact way.

Theorem 5 (Epstein et al., 2009) *For a single-source connection game Γ on a directed series-parallel graph a strong equilibrium always exists and can be computed in polynomial time.*

The existence result follows easily by observing that the Flow-LP for single-source games on directed series-parallel graphs has integrality gap 1. A proof can be derived from (Prodon et al., 1985). Solving this LP resembles the construction of Epstein et al. (2009).

Another class, in which the above LP can be used to show existence of SE, are MST games as mentioned in the introduction. MST games are single-source games, in which in every vertex of G is a sink node for at least one player.

Theorem 6 *In every MST game Γ there is a strong equilibrium, which can be computed in polynomial time.*

For the problem in directed graphs, a SE can be computed from dual solutions of the LP (Tamir, 1991). One of these dual solutions is the core solution derived for the original non-emptiness proof (Granot and Huberman, 1981). In this solution, each player k pays exactly for the unique arc of the tree leaving his sink t_k . This rule has also been described by Bird (1976). It requires an easy argument to see that it yields a SE, even for the MST game in undirected graphs.

While in these cases we have guaranteed existence and efficient algorithms to compute SE, the problem of deciding the existence of SE is NP-hard. This follows from a simple adjustment, which allows to interpret UFL games as single-source connection games on directed graphs.

Corollary 6 *It is NP-hard to decide if a given single-source connection game Γ on a directed graph has a strong equilibrium.*

4 Strong Equilibria beyond Linear Programming

4.1 Connection Games

For set cover and facility location games the integrality gap condition provides a complete characterization of games Δ having core solutions. With our theorems we obtain a complete characterization also for the existence of SE in strategic games Γ . For network design games like the connection game, the integrality gap condition is sufficient to show existence of SE and non-emptiness of the core, but it is not necessary. A tight characterization of games with non-empty core has not been obtained so far.

For strategic games and SE it has been shown by Epstein et al. (2009) that there is a single-source connection game without SE, but the corresponding cooperative game to their example does not allow a core solution as well. By Proposition 3, however, this is a prerequisite for SE existence. Coalitional connection games with an empty




Fig. 1 A single-source connection game with 3 players, a non-empty core, but without a SE. R^* is an MST of G and consists of all edges of cost 20.

core (and thus without SE) have already been presented by Granot and Maschler (1998). We here show that even a spanning property of the optimum solution R^* is not sufficient to guarantee SE existence or to obtain SE from core solutions. This implies that the relation between core and SE is not as robust as for the other games considered previously. A complete characterization of the existence of SE in (single-source) connection games remains as an open problem.

Lemma 1 *There are corresponding strategic and coalitional single-source connection games Γ and Δ such that R^* is a MST of G and Δ has a core solution but Γ has no strong equilibrium.*

Proof Our example game is shown in Fig. 1. It is based on a game presented by Granot and Maschler (1998), which consisted only of the three lower layers up to node s' . It was shown that this game has an empty core, but R^* passes through all vertices of G . This also implies that there can be no SE.

To obtain our game in Fig. 1, we added the new source s and an edge of cost 20 to the old source s' . Then the constraints for the contributions of the coalitions allow a feasible cost sharing by assigning each player a share of $160/3 \approx 53.33$. This removes the incentives to deviate on a global scale, which is sufficient for non-emptiness of the core. On a local scale, however, the instable structure up to s' is still intact. The additional contributions towards (s', s) do not change the strategic incentives within the lower parts of the graph. It can be verified that in this game no SE exists. This proves the lemma. \square

4.2 Terminal Backup Games

In this section we study games based on the terminal backup problem (Anshelevich and Caskurlu, 2011b; Anshelevich and Karagiozova, 2011). In this game there is a graph $G = (V, E)$, each player is a vertex ($K \subset V$), and resources are the edges with

costs $c(e) \geq 0$. Each player strives to be connected to at least $d - 1$ other player vertices, for $d \geq 2$. It has been shown by Anshelevich and Karagiozova (2011) that the terminal backup problem can be solved in polynomial time for $d = 2$. Here we show that every core solution can be turned into a SE for these games. In addition, we show how to decide if a game has a SE and how to obtain SE in polynomial time if they exist.

Theorem 7 *If a terminal backup game Δ with $d = 2$ has a non-empty core, then the strategic game Γ has a strong equilibrium.*

Proof Suppose there is a core solution in Δ , but there is no SE in Γ . We first adjust the graph such that R^* is only composed path components with two player vertices at the ends, or of star components with at most three player vertices which are the leaves of the star. This adjustment can be achieved as follows. Clearly, every component R^* is either a path or a star component, and every path end or star leaf is a player vertex. If there is a star component with a player vertex at the center, we can introduce an auxiliary vertex, make this the player vertex, and connect it to the star center with an edge of cost 0. For stars with more than three leaves, we can replace the non-player center vertex by a clique of sufficiently many vertices and clique edges of cost 0. In this way, we can split the star up into paths and at most one star of three player vertices. Note that these adjustments change the structure but not the cost of any solution to the underlying optimization problem.

We can allocate the cost shares from a core solution γ as follows. Consider a star component in the optimum solution R^* , which we can assume to consist of exactly three player vertices at the leaves. We will see that in a core solution γ , each player pays a cost of at most the connection to the center. Let the player vertices in the star be v_1, v_2 and v_3 and the star center w . We denote by $P(u_1, w)$ the path between u_1 and w in the star and by $c(u_1, w)$ the cost of $P(u_1, w)$ (for players 2 and 3 similarly). For contradiction, suppose that the core cost share of player 1 is $\gamma_1 \geq c(u_1, w) + \varepsilon$ with $\varepsilon > 0$. Then, as players 1 and 2 could deviate to the path $P(u_1, w) \cup P(u_2, w)$, the core constraints imply $\gamma_1 + \gamma_2 \leq c(u_1, w) + c(u_2, w)$. Hence, $\gamma_2 \leq c(u_2, w) - \varepsilon$. Note that the same argument holds for player 3, so $\gamma_3 \leq c(u_3, w) - \varepsilon$. Thus, in total γ_1, γ_2 and γ_3 do not pay completely for the star. The remaining cost of 2ε must hence be contributed by some other player not in the star component. However, the core constraints imply that no subset of players pays more than the cost of their component in R^* . This implies that whenever $\varepsilon > 0$, the cost of R^* cannot be fully paid for, a contradiction. Thus, the core cost share allows each player to pay completely for every edge of $P(u_i, w)$ in his star, and this is how we assign players to pay in their strategy.

If the component is a path, we allocate the cost shares such that each player k considers the edges of the path consecutively starting from his end vertex. He tries to pay them completely in this order until his core cost share γ_k is exhausted. Hence, there is at most one edge on the path for which the cost is shared by the two players. Because the core constraints forbid any subset of players to pay more than the cost of their component in R^* , the players pay for the cost of the path exactly if and only if γ is a core solution.

For the sake of contradiction assume that this allocation is not a SE. Suppose C is a violating coalition of players. In their improvement the players of C can improve

by changing their connections and create a new component. If this new component is paid for fully by the players in C , this corresponds to a constraint considered for the core solution. Hence, the players in such a new component cannot all profit from such a deviation.

On the other hand, suppose players in C use edges to create their new component, for which (part of) the cost is paid for by players not in C . Note that in our assignment, the edges that player k contributes to form a consecutive path starting at his terminal. He shares the cost of at most one edge with one other player at the end of this subpath. Hence, if in a deviation players of C use an edge fully paid for by $k \notin C$, we can include k into the deviating coalition, as sticking to his current strategy will guarantee that his connection requirement will be satisfied in some new component created by C as well. Furthermore, all players from the component of k that are not in C can be included into C , as they will all remain connected by sticking to their current strategies. If the deviation of C uses an edge for which the cost is shared, we can add both players that currently pay for the cost to C , because by sticking to their strategies they remain connected in the deviation as well.

Finally, we can redistribute the costs among all players of the enlarged coalition C such that everybody improves and pays a strictly smaller amount than before. This again results in a set of improving players that pays completely for their component. However, as such deviations are covered by the core constraints, this is a contradiction to the cost shares being a core solution. This completes the proof of the theorem. \square

The above property allows us to efficiently determine if SE and core solutions exist and to compute them in polynomial time if they exist.

Corollary 7 *There is a polynomial time algorithm to determine if a coalitional terminal backup game Δ with $d = 2$ has a core solution and if the strategic game Γ has a strong equilibrium. If they exist, a core solution and a strong equilibrium can be computed in polynomial time.*

Proof We can compute an optimal solution in polynomial time. We then decide if a core solution exists as follows. As outlined above, the structure of the problem allows to transform optimal solutions into compositions of components for two or three players. Thus, possible deviations from the grand coalition by coalitions of size 4 or larger can be reduced to collections of deviations by coalitions of two or three players. There are only a polynomial number of such coalitions, and the optimum solution for each such coalition can be found in polynomial time for each of them. Hence, the set of inequalities necessary to characterize the core is only of polynomial size and can be obtained in polynomial time. Thus, we can check in polynomial time if this set of inequalities has a solution and in this way obtain a member of the core. Given a core solution, we can use the computed optimum solution and our structural insight about SE to find the appropriate allocation of payments to edges in polynomial time. \square

For larger connectivity requirements of $d \geq 4$ we construct games where the consecutive payment condition of Theorem 7 is violated. In this case, a core solution cannot be turned into a SE.

Lemma 2 *For any $d \geq 4$ there is a coalitional terminal backup game Δ with a core solution and a corresponding strategic game Γ without a strong equilibrium.*

In fact, our example game can be derived directly with the single-source connection game in Fig. 1 above. We simply replace the source s by a clique of 4 or more terminals and 0-cost edges.

4.3 Network Cutting Games

In this section we briefly discuss a network cutting game, in which there is a graph $G = (V, E)$ and each player strives to disconnect a subset $U_k \subset V$ from another subset $V_k \subset V$. Each edge $e \in E$ has a cost $c(e) > 0$ for disconnection. This approach yields coalitional and strategic games based on a variety of minimum-cut problems like s - t -cut, multicut, multiway cut, etc. It was introduced and studied with respect to NE in the special cases of multiway cut and multicut by Anshelevich et al. (2010).

More formally, for each player k denote by \mathcal{P}_k the set of all paths in G from a node in U_k to a node in V_k . When we introduce a variable x_e for each edge $e \in E$, then for each path $P \in \mathcal{P}_k$ player k has the constraint $\sum_{e \in P} x_e \geq 1$. Note that these are simple 0/1-covering constraints, and thus the resulting integer program is a special case of the set cover integer program presented above. In particular, we can simply regard paths as elements and edges as sets. This implies that if the integrality gap is 1, we have existence of core solutions and SE. For instance, this holds on directed and undirected graphs for single-source games that have $U_k = \{u\}$ for each $k \in K$.

Theorem 8 *If the Covering-LP has an integral optimum solution, then the strategic network cutting game Γ has a strong equilibrium.*

Note that there is an important detail in this observation. While in the set cover game every element (i.e., every path) is a player, in the cutting game players strive to cover multiple elements (i.e., cut multiple paths). The previous theorem still holds, because by clustering elements we simply reduce the granularity of possible coalitions to those, which can be obtained by the union of sets \mathcal{P}_k . In fact, by this transformation we increase the set of games that allow a strong equilibrium and a core solution.

Proposition 4 *There are network cutting games Γ with strong equilibria, for which the underlying network cutting problem has an integrality gap of more than 1.*

Proof Consider a network multiway cut game, in which every player k has a vertex $u_k \in V$ and wants to disconnect it from every other player vertex, i.e., $V_k = \{u_j : k \neq j \in K\}$. Consider a star, in which the player vertices are exactly the leaves and all edges have cost 1. This class of instances is known to have the maximum integrality gap of $2 - 2/|K|$ for the covering LP of the network multicut problem. In particular, the fractional optimum solution assigns each edge to be in the cut with $x_e = 1/2$, while the integral optimum fully cuts all but one edge. In a SE we pick one player k to be *uncut*. Each other player $j \neq k$ is assigned to purchase the edge incident to s_j completely. Note that every coalition C without the uncut player k must pay at least




Fig. 2 A multicut game on a directed graph with 2 players and a non-empty core. The game has no NE.

$|C|$ to remain disconnected from k . Every other coalition C must pay at least $|C| - 1$. Hence, no coalition can reduce their payments in sum, and the existence of a SE and a core solution follows. \square

A similar observation can be made for multiway cut games, in which geometric LP relaxations (Calinescu et al., 2000) have an integrality gap of more than 1. In general network cutting games, however, the set of strategic games with a SE is not equivalent to the set of cooperative games with a non-empty core.

Lemma 3 *There are corresponding conditional and strategic network cutting games Δ and Γ such that Δ has a core solution but Γ has no strong equilibrium.*

Proof For undirected graphs we consider two players and a star graph. We set $U_1 = \{u_1\}$, $U_2 = \{u_2\}$, $V_1 = \{v_1\}$ and $V_2 = \{u_1, v_1\}$. The edge costs to the center node w are $c(u_1, w) = c(v_1, w) = 2$ and $c(u_2, w) = 3$. The set of core solutions is $\gamma_1 = 2 - \varepsilon$ and $\gamma_2 = 2 + \varepsilon$ for $\varepsilon \in [0, 1]$. Note that the unique optimum solution is to cut (u_1, w) and (v_1, w) . In such a solution, however, if $|s_1| > 0$, player 1 can unilaterally improve by removing the larger of his payments. Player 2 does not pay for both edges, because paying only for (u_2, w) is cheaper.

For directed graphs we can even leave $V_2 = \{v_2\}$ as a singleton. We transform the graph to the one shown in Fig. 2. A similar argument shows non-existence of SE. In particular, none of the edges of cost 100 is cut by the players. In the optimum, the two edges of cost 2 are cut. However, player 1 requires only one of them to cut his path. Hence, if $|s_1| > 0$, player 1 will unilaterally deviate and drop the larger of his contributions. Thus, player 2 would have to pay fully for both edges, but for him cutting the edge of cost 3 is cheaper. \square

This construction implies that when we relax the assumption that *every* element or terminal is a player in a set cover or facility location game, equivalence between core and SE does not hold anymore. On another note, the proof shows absence of NE in general strategic network cutting games on undirected games. For directed graphs the absence of NE holds true even for minimum multicut games, in which U_k and V_k are singleton sets for all players $k \in K$.

5 Extensions

5.1 Fractional and Non-Binary Resources

Apart from equivalence of core and SE, a natural question is to characterize cases when we can derive SE from core solutions using linear programming. This was possible when the integrality gap was 1 in all the cases described above. With exception of the CCRFL games, all games studied above yield linear constraints that fall into one of two classes. One type of constraint is $\sum_i x_i \geq 1$, i.e., a simple covering constraint with 0/1 coefficients, by which we can express exactly the vertex and set cover games. The other constraint type is $y_i - \sum_j x_{ij} \geq 0$, i.e., a coordination constraint that requires a resource to become bought when at least one player uses it. This second type of constraint allows us to treat facility location and network design cut games. What happens if we slightly generalize these constraints?

As an example let us first consider dropping the integrality requirement. Using results on Owen's linear production model one can show, for instance, that vertex cover games always allow a core solution if vertices and sets can be bought in *fractional* amounts. Does a SE also exist for strategic games in these cases? To answer this question we must adjust the strategic game to allow vertices to be bought fractionally. The obvious adjustment is to assign a fraction proportional to the total payment. In a state s of the strategic *fractional vertex cover game* a vertex v is bought to the degree $x_v = \sum_{k \in K} s_k(e)/c(v)$. For a player k corresponding to edge $e = (u, v)$ the individual cost is $|s_k|$ if $x_u + x_v \geq 1$ and prohibitively large otherwise.

A second, closely related variant is the case when we keep the integrality condition, but we increase the covering requirements and allow multiple units of a resource to be bought. In particular, we change the constraints to a type $\sum_i x_i \geq b$, where $b > 0$ and $x_i \in \mathbb{N}$. As for the fractional games the total payments of the players determine the number of units bought of a resource. We term these games *non-binary vertex cover games*. More formally, in a state s we have $x_u = \lfloor \sum_{k \in K} s_k(u) \rfloor$. Player k corresponding to edge (u, v) has a required coverage of $b_k \in \mathbb{N}$ and individual cost $c_k(s) = |s_k|$ if $x_u + x_v \geq b_k$ and prohibitively large otherwise.

Note that for both of these game classes Propositions 3 and 2 continue to hold. In contrast to our results above, however, we show next that there might be no SE – although non-emptiness of the core can be established via the same linear programming machinery that was used before.

Theorem 9 *There are corresponding strategic and coalitional fractional or non-binary vertex cover games Δ and Γ such that Δ has a core solution but Γ has no strong equilibrium.*

Proof For both variants the proof follows with a triangle, vertex costs $c(u) = 3$, $c(v) = 5$, and $c(w) = 7$, and players 1 to 3 corresponding to edges (u, w) , (u, v) and (v, w) , respectively.

In the fractional game the unique optimum solution to the underlying vertex cover problem is $x_u^* = x_v^* = x_w^* = 1/2$, and the unique core solution is $\gamma_1 = 2.5$, $\gamma_2 = 0.5$ and $\gamma_3 = 4.5$. Proposition 2 yields that x^* has to be purchased in every SE, but no player is willing to contribute to w . We obviously must have $s_2(w) = 0$. If $s_1(w) > 0$, player 1

can deviate unilaterally and achieve the amount $s_1(w)/7$ of coverage by contribution to u with less payments. The same holds for player 2 and vertex v .

For the non-binary version, we set all covering requirements to $b_1 = b_2 = b_3 = 4$. Then the unique optimum x^* to the underlying vertex cover problem and the unique core payments γ are the same as before scaled by factor 4. Observe that we have an integrality gap of 1 in this game. The core solution is unique, so with Proposition 3 we know that in every SE $|s_1| = 10$ and $|s_2| = 2$. This implies $4 \leq s_1(w) \leq 6$. By removing this payment from w , player 1 reduces the number of units bought of w by exactly 1. However, he can obtain an additional unit of u at a cost of 3. This yields a profitable unilateral deviation and proves the theorem. \square

This shows that in the class of non-binary vertex cover games neither non-emptiness of the core nor an integrality gap of 1 can guarantee the existence of SE.

5.2 Approximate Equilibria

We have presented a method to derive SE in strategic cost sharing games via linear programming. A disadvantage of the concept of SE is that they might not exist in a game. However, our approach proves to be applicable even to approximate SE. Using primal-dual algorithms we can compute (α, β) -SE with small (constant) ratios in polynomial time for vertex cover, set cover, and facility location games. The proof for the following theorem can be derived directly from arguments in (Cardinal and Hofer, 2010).

Theorem 10 *There are efficient primal-dual algorithms to compute $(2, 2)$ -SE for vertex cover, (f, f) -SE for set cover (where f is the maximum frequency of any element in the sets), and $(3, 3)$ -SE for metric UFL games in polynomial time.*

Proof The proof follows with a close observation of the results in (Cardinal and Hofer, 2010). In these works, we have observed that the results stated in the theorem hold for (α, β) -approximate Nash equilibria with the same ratios in vertex cover, set cover, and metric UFL games, even for games in which a single player has control of more than one edge, element or terminal, respectively.

To outline the general idea of the proof, consider the case of $(2, 2)$ -NE in vertex cover games studied by Cardinal and Hofer (2010). The primal-dual algorithm makes a single iteration through all the edges in arbitrary order. For a chosen edge, it raises payments at both endvertices until the total contribution to one the vertices suffices to pay the cost. This is done until all edges are covered, and then the algorithm terminates. If players own multiple edges, their total payments are made up by the sum of payments assigned to their single edges. Obviously, the payments of single edges assigned by the algorithm are independent of which player owns which edge. Furthermore, observe that a deviation of a player owning multiple edges is equivalent to a coordinated deviation by the coalition of single edge players. Thus, the proof that the algorithm computes $(2, 2)$ -NE shows that the state computed by the algorithm is a $(2, 2)$ -NE independent of how the edges are owned by the players. Hence, no subset of edges has a deviation that decreases their payments in sum by a factor of strictly

more than 2. In this case we obviously cannot have a deviation in which *every* player of a coalition reduces his cost by a factor of strictly more than 2. These observations yield the result for vertex cover games.

The main properties are that (1) the result holds for Nash equilibria even when players can own multiple edges, (2) there is a cost per element and the cost for a player is the sum of his element costs, and (3) the algorithm assigns the element cost independently of which player owns the element. These conditions hold also for the algorithms presented for set cover and facility location games (Cardinal and Hoefer, 2010) and hence yield the theorem. \square

6 Conclusions and Open Problems

We have studied cost sharing in strategic and cooperative games and shown some interesting connections between coalitional stability concepts. In simple games such as vertex cover, set cover, facility location, MST, or simple terminal backup games, existence of core and SE is equivalent. In these games, even algorithms for computation of core solutions can be used to compute SE. Here it seems that the cooperative framework is an appropriate abstraction as the consideration of strategic incentives does not lead to significantly different properties. In more general games, however, the differences between core and SE highlight the fact that strategic incentives have a non-trivial effect on stability and fairness in a cost sharing scenario. More work needs to be done to fully understand and distinguish these effects.

There are a number of open problems that stem from our work. In (single-source) connection, network cutting and fractional and non-binary games the use of linear programming duality does not necessarily yield a complete characterization of the games that admit SE. In these games and other interesting variants of cost sharing in network design our work opens up numerous interesting research problems regarding the characterization and computation of exact and approximate SE.

More generally, we believe that the linkage between core and strong equilibrium could be present in other cost sharing games, which go beyond the classes of games treated in this paper. Exploring these classes of games is an interesting avenue for further research. More concretely, our games have linear programming formulations that lie within Owens linear production model. Non-emptiness of the core, however, can also be shown within a more general class of problems. This more general framework, termed generalized linear production model by Granot (1986), has a non-additive structure, and it encompasses for instance the cut-based LP-formulation for Steiner Network problems (Skorin-Karpov, 1995). It is a fascinating open problem to see if this framework can also be used to derive exact and approximate SE in strategic cost sharing games.

Acknowledgements The author would like to thank Elliot Anshelevich and Bugra Caskurlu for valuable discussions and feedback about the results in this paper.

References

- Susanne Albers. On the value of coordination in network design. *SIAM J. Comput.*, 38(6):2273–2302, 2009.
- Nir Andelman, Michal Feldman, and Yishay Mansour. Strong price of anarchy. *Games Econom. Behav.*, 65(2):289–317, 2009.
- Elliot Anshelevich and Bugra Caskurlu. Price of stability in survivable network design. *Theory Comput. Syst.*, 49(1):98–138, 2011a.
- Elliot Anshelevich and Bugra Caskurlu. Exact and approximate equilibria for optimal group network formation. *Theoret. Comput. Sci.*, 412(39):5298–5314, 2011b.
- Elliot Anshelevich and Adriana Karagiozova. Terminal backup, 3D matching, and covering cubic graphs. *SIAM J. Comput.*, 40(3):678–708, 2011.
- 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, 2008a.
- Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4:77–109, 2008b.
- Elliot Anshelevich, Bugra Caskurlu, and Ameya Hate. Strategic multiway cut and multicut games. In *Proc. 8th Intl. Workshop Approximation and Online Algorithms (WAOA)*, pages 1–12, 2010.
- Robert Aumann. Acceptable points in general cooperative n-person games. In *Contributions to the Theory of Games IV*, volume 40 of *Annals of Mathematics Study*, pages 287–324. Princeton University Press, 1959.
- C. Bird. On cost allocation for a spanning tree: A game theoretic approach. *Networks*, 6:335–350, 1976.
- Gruia Calinescu, Howard Karloff, and Yuval Rabani. An improved approximation algorithm for multiway cut. *J. Comput. Syst. Sci.*, 60(3):564–574, 2000.
- Jean Cardinal and Martin Hoefer. Non-cooperative facility location and covering games. *Theoret. Comput. Sci.*, 411(16–18):1855–1876, 2010.
- Ho-Lin Chen, Tim Roughgarden, and Gregory Valiant. Designing network protocols for good equilibria. *SIAM J. Comput.*, 39(5):1799–1832, 2010.
- Xiaotie Deng, Toshihide Ibaraki, and Hiroshi Nagamochi. Algorithmic aspects of the core of combinatorial optimization games. *Math. Oper. Res.*, 24(3):751–766, 1999.
- Amir Epstein, Michal Feldman, and Yishay Mansour. Strong equilibrium in cost sharing connection games. *Games Econom. Behav.*, 67(1):51–68, 2009.
- Ulrich Faigle, Sándor Fekete, Winfried Hochstättler, and Walter Kern. On approximately fair cost allocation in Euclidean TSP games. *OR Spektrum*, 20(1):29–37, 1998.
- Michel Goemans and Martin Skutella. Cooperative facility location games. *J. Algorithms*, 50(2):194–214, 2004.
- Michel Goemans and David Williamson. A general approximation technique for constrained forest problems. *SIAM J. Comput.*, 24(2):296–317, 1995.
- Daniel Granot. A generalized linear production model: A unifying model. *Math. Prog.*, 34:212–222, 1986.

- Daniel Granot and Gur Huberman. On minimum cost spanning tree games. *Math. Prog.*, 21:1–18, 1981.
- Daniel Granot and Michael Maschler. Spanning network games. *Intl. J. Game Theory*, 27:467–500, 1998.
- Martin Hoefer. Non-cooperative tree creation. *Algorithmica*, 53(1):104–131, 2009.
- Martin Hoefer. Strategic cooperation in cost sharing games. In *Proc. 6th Intl. Workshop Internet & Network Economics (WINE)*, pages 258–269, 2010.
- Martin Hoefer. Competitive cost sharing with economies of scale. *Algorithmica*, 60(4):743–765, 2011.
- Martin Hoefer and Piotr Krysta. Geometric network design with selfish agents. In *Proc. 11th Conf. Computing and Combinatorics (COCOON)*, pages 167–178, 2005.
- Nicole Immorlica, Mohammad Mahdian, and Vahab Mirrokni. Limitations of cross-monotonic cost sharing schemes. *ACM Trans. Algorithms*, 4(2), 2008. Special Issue SODA 2005.
- Kamal Jain and Mohammad Mahdian. Cost sharing. In Noam Nisan, Éva Tardos, Tim Roughgarden, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 15. Cambridge University Press, 2007.
- Kamal Jain and Vijay Vazirani. Applications of approximation algorithms to cooperative games. In *Proc. 33rd Symp. Theory of Computing (STOC)*, pages 364–372, 2001.
- 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.
- Stefano Leonardi and Piotr Sankowski. Network formation games with local coalitions. In *Proc. 26th Symp. Principles of Distributed Computing (PODC)*, pages 299–305, 2007.
- Nimrod Megiddo. Cost allocation for Steiner trees. *Networks*, 8(1):1–6, 1978.
- Guillermo Owen. On the core of linear production games. *Math. Prog.*, 9:358–370, 1975.
- Martin Pál and Éva Tardos. Group strategyproof mechanisms via primal-dual algorithms. In *Proc. 44th Symp. Foundations of Computer Science (FOCS)*, pages 584–593, 2003.
- A. Prodon, T.M. Libeling, and H. Gröflin. Steiner’s problem on two-trees. Technical report, Département de Mathématiques, EPF Lausanne, 1985. Working paper RO 850315.
- Darko Skorin-Karpov. On the core of the minimum cost Steiner tree game in networks. *Annals of Operations Research*, 57:233–249, 1995.
- Arie Tamir. On the core of network synthesis games. *Math. Prog.*, 50:123–135, 1991.
- Arie Tamir. On the core of cost allocation games defined on location problems. *Transportation Sci.*, 27:81–86, 1993.
- Richard Wong. A dual ascent approach for Steiner tree problems on a directed graph. *Math. Prog.*, 28(3):271–287, 1984.
- H. Peyton Young. Cost allocation. In Robert Aumann and Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 2, chapter 34, pages 1194–1235. Elsevier/North-Holland Science Publishers, 1994.

A Connection-Restricted Facility Location

In a *CRFL problem* there is a set T of terminals and a set F of facilities. We set $n_t = |T|$ and $n_f = |F|$. In addition to the UFL problem each facility has a set of allowable subsets $\mathcal{A}_f \subseteq 2^T$. The goal is to open a subset of facilities and buy a set of connections of minimum total cost, such that each terminal is connected to an opened facility, and the set of terminals connected to each opened facility f is in \mathcal{A}_f . In the *CRFL game* each player owns a terminal, i.e., $K = T$. The constraint of player t is satisfied if there is a bought connection (t, f) to some opened facility f , and the subset of terminals that have a bought connections to f is in \mathcal{A}_f . We can formalize the CRFL problem by an integer program as follows:

$$\begin{aligned} \text{Min} \quad & \sum_{f \in F} c(f)y_f + \sum_{t \in T} c(t, f)x_{tf} \\ \text{subject to} \quad & \sum_{f \in F} x_{tf} \geq 1 && \text{for all } t \in T \\ & (y_f, x_{1f}, \dots, x_{n_f f}) \in \mathcal{A}_f && \text{for all } f \in F \\ & y_f, x_{tf} \in \{0, 1\} && \text{for all } t \in T, f \in F, \end{aligned}$$

where

$\mathcal{A}_f = \{(0, \dots, 0)\} \cup \{(1, \chi_{A_f}) \mid A_f \subseteq T \text{ feasible for } f\} \subseteq \{0, 1\}^{n_t+1}$, and χ_{A_f} denotes the characteristic vector of the subset A_f .

We here concentrate on a subclass of *closed games* (denoted CCRFL). In these games the sets \mathcal{A}_f are downward closed, i.e., for every $A \subseteq A' \in \mathcal{A}_f$ we have $A \in \mathcal{A}_f$. Note that this class encompasses a large variety of facility location problems considered in the literature, e.g., with capacity or incompatibility constraints.

Theorem 11 *If a CCRFL game Δ has a non-empty core, then the strategic game Γ has a strong equilibrium.*

Proof Following the argumentation in (Goemans and Skutella, 2004) it is possible to use the conic hull of the sets \mathcal{A}_f to derive a linear relaxation:

$$\begin{aligned} \text{Min} \quad & \sum_{f \in F} c(f)y_f + \sum_{t \in T} c(t, f)x_{tf} \\ \text{subject to} \quad & \sum_{f \in F} x_{tf} \geq 1 && \text{for all } t \in T \\ & (y_f, x_{1f}, \dots, x_{n_f f}) \in \text{cone}(\mathcal{A}_f) && \text{for all } f \in F. \end{aligned}$$

For this program a dual can be given by

$$\begin{aligned} \text{Max} \quad & \sum_{t \in T} \gamma_t \\ \text{subject to} \quad & \sum_{t \in A_f} \gamma_t \leq c(f) + \sum_{t \in A_f} c(t, f) \\ & \text{for } f \in F \text{ and } A_f \in \mathcal{A}_f. \end{aligned}$$

Now we can apply similar arguments as before. An integral optimum solution (x^*, y^*) to the LP-relaxation represents a partition of the terminal set T into a collection of feasible sets A_f^* , one for each facility f . The constraints corresponding to these sets hold with tightness, and we can assign each player t to pay for her terminal the amount $s_t(t, f) = \min\{\gamma_t^*, c(t, f)\}$ as connection cost to f with t connected to f , in which γ^* is the optimum solution to the dual. For the opening costs $s_t(f) = \max\{\min\{c(f), \gamma_t^* - c(t, f)\}, 0\}$. Note that such an assignment is always possible due to \mathcal{A}_f being downward closed. In particular, no player t is required to pay for the connection cost of any other player. Thus, no coalition of players can improve by simply dropping payments.

In total this pays exactly for all costs of the solution by duality. Suppose there is a violating coalition C . This coalition must be able to connect their terminals differently at a cheaper total cost. Consider the strategy vector after the coalition has changed its strategy. Each member $t' \in C$ must again be part of some $A_{f'}$ for some facility f' , for which the total (connection + opening) costs are fully paid for. In particular, the new payments exactly pay for $c(f') + \sum_{t \in A_{f'}} c(t, f')$. Note that no player has increased his payments, but t' has strictly decreased his payments. This means that the original payments coming from γ^* violate the dual constraint corresponding to $\mathcal{A}_{f'}$. This is a contradiction to γ^* being the optimal dual solution. \square

The recognition of SE can be done similarly as for UFL games.

Corollary 8 *Given a strategy vector for a CCRFL game Γ we can verify in polynomial time if it is a SE.*